Volume 7, Number 2 June 2011

IJCB INTERNATIONAL JOURNAL OF CENTRAL BANKING

INTERNATIONAL JOURNAL OF CENTRAL BANKING
Volume 7, Number 2 June 2011

IJCB International Journal of Central Banking

Volume 7, Number 2 June 2011

*The contents of this journal, together with additional materials provided by article authors, are available without charge at **www.ijcb.org**.*

The *International Journal of Central Banking* is published quarterly (ISSN: 1815-4654). Online access to the publication is available free of charge at **www.ijcb.org**. Individual print subscriptions are available at an annual rate of $100 (USD).

Print subscription orders may be placed online at www.ijcb.org, by phone (001 415 974 2035), via fax (001 415 974 2168), or by e-mail (editor@ijcb.org).

Renewals, claims, address changes, and requests for permission to reprint material from this journal should be addressed to:

International Journal of Central Banking
Economic Research Department
Federal Reserve Bank of San Francisco
101 Market Street
San Francisco, CA 94105
USA

Phone: 001 415 974 2035
Fax: 001 415 974 2168
E-mail: editor@ijcb.org

ISSN: 1815-4654

Imperfect Central Bank Communication: Information versus Distraction*

Spencer Dale,[a] Athanasios Orphanides,[b] and Pär Österholm[c]
[a]Bank of England
[b]Central Bank of Cyprus
[c]Sveriges Riksbank

Much of the information communicated by central banks is noisy or imperfect. This paper considers the potential benefits and limitations of central bank communications in a model of imperfect knowledge and learning. It is shown that the value of communicating imperfect information is ambiguous. If the public is able to assess accurately the quality of the imperfect information communicated by a central bank, such communication can inform and improve the public's decisions and expectations. But if not, communicating imperfect information has the potential to mislead and distract. The risk that imperfect communication may detract from the public's understanding should be considered in the context of a central bank's communications strategy. The risk of distraction means the central bank may prefer to focus its communication policies on the information it knows most about. Indeed, conveying more certain information may improve the public's understanding to the extent that it "crowds out" a role for communicating imperfect information.

JEL Codes: E52, E58.

*The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Bank of England, the Central Bank of Cyprus, or the Sveriges Riksbank. We would like to thank the editor Andy Levin, three anonymous referees, Helge Berger, Marcel Fratzscher, Lars Svensson, Tim Taylor, Daniele Terlizzese, John Williams, Tony Yates, and seminar participants at the Bank of England, Uppsala University, the Banca d'Italia/CEPR Conference on Money, Banking and Finance in Rome, the 2008 AEA Annual Meeting in New Orleans, and the Swiss National Bank Research Conference Alternative Models for Monetary Policy Analysis in Zürich for helpful comments. Österholm gratefully acknowledges financial support from the Jan Wallander and Tom Hedelius Foundation. Corresponding author: Pär Österholm, Sveriges Riksbank, 103 37 Stockholm, Sweden. E-mail: par.osterholm@riksbank.se. Tel: +46 8 787 0162.

1. Introduction

Over the past two decades, central banks' approach to communications has undergone a sea change. The cultivation of secrecy and mystique has been replaced by a zeal for openness and transparency. Although the benefits of an open and transparent monetary policy process are now widely recognized and understood, most central banks are still grappling with exactly how best to achieve that aim. An approach of "more information is always better" is neither sufficient nor correct. There are costs as well as benefits associated with communicating ever-increasing amounts of noisy and complex information. As central banks seek to continue their progress toward greater openness, these costs and benefits—and their implications for the design of central bank communications—need to be better understood.

This paper considers the potential benefits and limitations of central bank communications in a model that emphasizes two important features of the practical environment in which central banks operate. First, it recognizes that monetary policy is conducted against a backdrop of imperfect knowledge: Central banks and the public have imperfect knowledge about the structure of the economy and about the shocks affecting the economy. Second, it argues that central bank communications should be thought about in terms of communication strategies: All central banks publish some information—the issue faced by central banks is how to combine various types of information in a way that best informs the public. These two features have important implications for the analysis of central bank communications.

That monetary policy is conducted in an environment of imperfect knowledge is central to understanding the potential benefits and limitations of central bank communications. Foremost, it helps to motivate the importance that central banks place on communication policies. As noted by Orphanides and Williams (2006), central bank communications have little role to play in models of rational expectations with perfect knowledge.[1] The recognition that the

[1]In a similar vein, Ben Bernanke, Chairman of the Federal Open Market Committee, recently argued that "Notably, in a world with rational expectations and in which private agents are assumed already to understand all aspects of the economic environment, talking about the effects of central bank communication would not be sensible, whereas models with learning accommodate the analysis of communication-related issues quite well. . . . In sum, many of the most interesting issues in contemporary monetary theory require an analytical framework that involves learning by private agents and possibly the central bank as well" (Bernanke 2007).

private sector has imperfect knowledge rationalizes the role central bank communications can play in aiding private-sector learning. But the limits on central banks' knowledge and understanding imply that much of the information communicated by central banks is uncertain or imperfect, such as economic forecasts or short-term policy guidance. This raises important issues concerning the public's ability to process and utilize imperfect, noisy information. Interpreted correctly, the communication of imperfect information can help to inform and improve the public's understanding and expectations. But if the public are not able to assess accurately the quality of the information conveyed by central banks, imperfect central bank communication has the potential to distract and result in worse economic outcomes.[2] Imperfect central bank communication is a doubled-edged sword that should be used with care.

This risk of distraction is most apparent when considered in the context of a central bank's communications strategy. Central banks have a wide range of information at their disposal. All central banks communicate some information; the challenge faced by central banks is to decide which information to publish—and in what form—to best aid the efficient functioning of the economy. An assessment of whether or not to publish a particular piece of information has to be made in the context of the information that is already being communicated. Importantly, the different types of information available to a central bank vary considerably in terms of their precision: Information about a central bank's inflation objective or the outcome of a policy meeting is more precise—less imperfect—than a central bank's economic forecast or guidance about the future path of policy. The risk that the public may be unable to assess correctly the quality of the information being conveyed means that central banks may prefer to communicate more certain information. Indeed, the disclosure of more certain information may "crowd out" a role for communicating imperfect information. That is, it may improve the precision of the public's understanding to a point at which it is no longer beneficial for the central bank to run the risk of communicating imperfect information.

---

[2]Similar concerns have been expressed by, for example, Issing (1999, 2000), Winkler (2000), Mishkin (2004), Macklem (2005), Woodford (2005), King (2006), and Sibert (2009).

This paper considers the role of central bank communication in an environment of imperfect knowledge and learning. Both the central bank and the private sector are assumed to have imperfect knowledge about the structure of the model and are engaged in perpetual learning. The risk of distraction is explored by considering the effect of publishing central bank forecasts on the accuracy of private-sector inflation expectations. Since both the central bank and the private sector are assumed to have imperfect knowledge of the economy, the central bank's forecasts are not necessarily more accurate than those of the private sector. If the private sector is able to assess correctly the precision of the central bank's forecasts, publishing the forecasts improves the accuracy of private-sector expectations. But if the private sector is not able to assess the quality of the central bank's forecasts, the value of publishing central bank forecasts is shown to be ambiguous. In particular, the private sector may inadvertently place too much weight on the central bank's forecast and so detract from the accuracy of its expectations.

Our results are qualitatively similar to those of Morris and Shin (2002), who show that the welfare effects of increased public disclosures are ambiguous.[3] But the mechanism underlying the Morris and Shin result is different from that considered here. Morris and Shin's analysis is based around a model of information heterogeneity in which agents are assumed to have a coordination motive arising from a strategic complementarity in their actions. The role that public information serves as a focal point for beliefs leads agents to attach excessive weight to such information. Thus, as the relative precision of the public signal deteriorates, the provision of public information can be detrimental to welfare. In the model considered in this paper, the possibility that the private sector may attach the incorrect weight to central bank information stems from their imperfect ability to assess the quality of such information: Agents may attach too much or too little weight to the information. As the relative precision of the central bank information declines, the costs associated with attaching too much weight to the information increase, such that there is growing risk that central bank communications

[3]In particular, they argue that "the better informed is the private sector, the higher is the hurdle rate of precision of public information that would make it welfare enhancing" (Morris and Shin 2002, p. 1529).

may act as a source of distraction. The private sector's ability to assess the quality of the central bank's information is hence important in the model used in this paper. If this ability is limited, there is a higher probability that communication of noisy central bank information will mislead the private sector.

The importance of viewing central bank communications as a strategy comprising numerous pieces of information is illustrated by considering a case in which the central bank is able to announce its inflation objective as well as publish its economic forecasts. It is shown that announcing an inflation target has the potential to "crowd out" a role for publishing economic forecasts.

The remainder of this paper is organized as follows: Section 2 outlines the model and informational assumptions, section 3 presents the main results, and section 4 concludes.

2. The Model

The model used to explore the effects of central bank communication is highly stylized. Both the central bank and the private sector produce forecasts of inflation. The central bank is assumed to use a structural model as the basis of its forecasting model, whereas the private sector relies solely on a reduced-form forecasting model. The choice of models reflects the dominant forecasting strategies used by central banks and private-sector forecasters, respectively, in many countries. In particular, the greater value that central banks tend to place on understanding and explaining the "economics" underlying their forecasts means that they often make greater use of structural economic models.[4] But the assumptions concerning the use and choice of different forecasting models are not important for what follows. All that matters is that the forecasts produced by the central bank and the private sector are distinct, and that there is a possibility that the central bank's inflation forecasts may be less accurate than those of the private sector.[5]

[4]See, for example, Harrison et al. (2005) for a discussion of the objectives underlying the design of the Bank of England's forecasting model.

[5]Similar results could be obtained, for example, by assuming that the central bank and the private sector used identical forecasting models but received different signals concerning the state of the economy.

Importantly, the models employed by the central bank and private sector are properly specified in the sense that they nest the correct structure of the economy and equilibrium dynamics that would prevail under a rational expectations equilibrium with perfect knowledge. The central bank and private sector update their model coefficients recursively using constant-gain least squares. This estimation algorithm is equivalent to applying weighted least squares where the weights decline geometrically with time.[6] As discussed by Orphanides and Williams (2004, 2006), the use of constant-gain learning—which has the property that learning is a never-ending (perpetual) process—can be justified by the central bank and private sector allowing for the possibility of structural change and therefore placing less weight on older data. The central bank and private sector recursively update their forecasting models each period and use their most recent estimates to generate inflation forecasts.

The effectiveness of central bank communication is evaluated according to its ability to improve the accuracy of private-sector inflation expectations. If the central bank publishes its inflation forecast, the private sector combines forecasts from its own model with those of the central bank in order to form expectations of future inflation. The weight attached to the central bank's forecast in this combination is determined by the historical forecasting performance of the central bank relative to that of the private sector's model, where importantly that weight is also recursively updated using constant-gain learning. That is, the private sector is perpetually learning about the (relative) quality of the central bank's forecasts. If the central bank does not publish its inflation forecast, the private sector's inflation expectations are based solely on its own model forecasts.

The model employed is similar to that used by Orphanides and Williams (2004) but has been extended in several dimensions. Inflation is determined according to a modified Lucas supply function as

$$\pi_t = \phi\pi^e_{t|t-1} + (1-\phi)\pi_{t-1} + \alpha x_t + e_t, \tag{1}$$

[6]Sargent (1993, 1999) and Evans and Honkapohja (2001) discuss properties of constant-gain learning.

Figure 1. The Timing of the Model

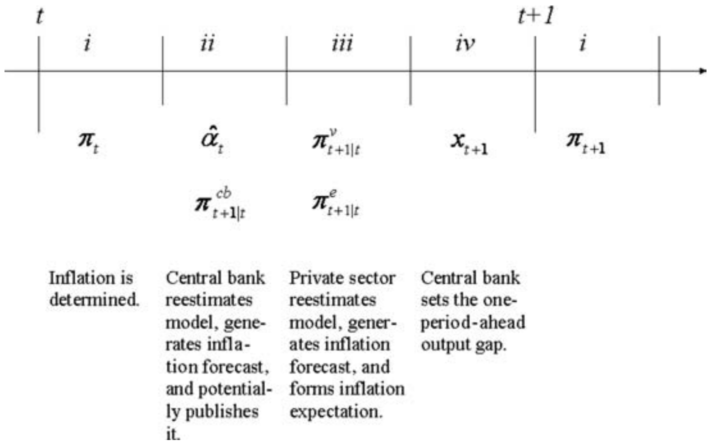

where π_t is inflation, $\pi^e_{t|t-1}$ is the private-sector expectation of time t inflation formed at $t-1$, x_t is the output gap, and e_t is a disturbance with properties $e_t \sim NID(0, \sigma_e^2)$. The output gap for period t is determined by the central bank in period $t-1$ according to its reaction function as

$$x_t = -\theta(\pi_{t-1} - \pi^*), \tag{2}$$

where π^* is the inflation target of the central bank.

Before considering the forecasting models used by the central bank and private sector, it is useful to clarify the timing of the model. As illustrated in figure 1, each time period is separated into four subphases. In phase (i), inflation in period t is determined and observed by the central bank and private sector. In phase (ii), the central bank reestimates its model, produces its forecast for inflation in period $t+1$, and decides whether to publish it. The private sector reestimates its forecasting model in phase (iii), uses the updated model estimates to generate its inflation forecast, and forms its expectation for inflation in period $t+1$ based on its own inflation forecast and (if published) the central bank's forecast. The private sector's inflation expectations for the next period are observed by the central bank. Finally, in phase (iv), the central bank decides its policy setting (the output gap) for the next period.

The central bank estimates a structural econometric model of the economy, namely the supply function in equation (1). To capture the possibility that central banks have imperfect knowledge of the economy, the central bank is assumed to know the correct form of the supply function and the true value of ϕ, but *not* the value of α.[7] The central bank therefore recursively estimates α each period using constant-gain least squares according to

$$\pi_t - \phi\pi^e_{t|t-1} - (1-\phi)\pi_{t-1} = \alpha_t x_t + \psi_t. \tag{3}$$

The regression coefficient α_t can be written as

$$\hat{\alpha}_t = \hat{\alpha}_{t-1} + \kappa^{cb}\left(R^{cb}_t\right)^{-1} x_t\left(\pi_t - \phi\pi^e_{t|t-1} - (1-\phi)\pi_{t-1} - \hat{\alpha}_{t-1}x_t\right), \tag{4}$$

where

$$R^{cb}_t = R^{cb}_{t-1} + \kappa^{cb}\left(x^2_t - R^{cb}_{t-1}\right) \tag{5}$$

and κ^{cb} is the gain.[8]

The central bank's inflation forecast is given by

$$\pi^{cb}_{t+1|t} = \frac{\hat{\alpha}_t\theta}{1-\phi}\pi^*_t + \frac{1-\phi-\hat{\alpha}_t\theta}{1-\phi}\pi_t, \tag{6}$$

where this forecast assumes that private-sector inflation expectations are formed according to the central bank's estimate of the reduced-form relationship that would be the solution under full information and rational expectations.[9]

[7]The assumption that the central bank knows ϕ but not α is not particularly realistic but allows us to introduce imperfect knowledge of the central bank in a "minimalistic" way, which is beneficial from a pedagogical viewpoint. Changing several aspects of the model at once can be obscuring, as it might be unclear what the relative importance of the different aspects is.

[8]This learning algorithm is standard in the literature; see, for example, Evans and Honkapohja (2001).

[9]This assumption is made for simplicity. In principle, the central bank could, for example, estimate a separate forecasting model for private-sector expectations and use this together with the structural model as the basis for its inflation forecasts. This would not affect the main qualitative results discussed here.

The private sector generates inflation forecasts using the AR(1) model

$$\pi_t = c_{0,t} + c_{1,t}\pi_{t-1} + v_t. \tag{7}$$

Just like the central bank, the private sector recursively updates its model estimates each period using constant-gain least squares, and we can express the regression coefficients $\hat{c}_t = (\hat{c}_{0,t}, \hat{c}_{1,t})'$ as

$$\hat{c}_t = \hat{c}_{t-1} + \kappa^v R_t^{-1} X_t(\pi_t - X_t'\hat{c}_{t-1}), \tag{8}$$

where

$$R_t = R_{t-1} + \kappa^v \left(X_t X_t' - R_{t-1}\right), \tag{9}$$

κ^v is the gain, and $X_t = (1, \pi_t)'$.

The model in equation (7) is then used to generate a model-based forecast for inflation according to

$$\pi^v_{t+1|t} = \hat{c}_{0,t} + \hat{c}_{1,t}\pi_t. \tag{10}$$

If the central bank does not publish its inflation forecast, the private sector's inflation expectation is simply formed as $\pi^e_{t+1|t} = \pi^v_{t+1|t}$. In general, though, we let the private sector's expectation for inflation be based on both its own inflation forecast and the forecast published by the central bank. Given the evidence in the literature suggesting that the central bank is a very good forecaster—see, for example, Romer and Romer (2000)—it seems natural that the private sector wants to take the central bank's forecast into account when forming its expectation. In particular, the private sector combines its own least-squares forecast in equation (10) with the central bank's according to

$$\pi^e_{t+1|t} = \gamma_t \pi^v_{t+1|t} + (1 - \gamma_t)\pi^{cb}_{t+1|t}. \tag{11}$$

In line with the principle of "optimal weights" suggested by Granger and Ramanathan (1984), the weight γ_t in this forecast combination is determined by the relative historical forecasting performances of the private sector and central bank. However, given that the private sector is assumed to be perpetually learning about the structure of the economy, rather than using fixed weights for the forecast combination, the private sector updates the weight each period as suggested by Diebold and Pauly (1987). Assuming that both the private sector and the central bank are generating unbiased forecasts, the private sector establishes γ_t by running the regression

$$\pi_t - \pi^{cb}_{t|t-1} = g_t\left(\pi^{v}_{t|t-1} - \pi^{cb}_{t|t-1}\right) + \chi_t \tag{12}$$

using constant-gain least squares; the gain in this procedure is denoted κ^f. For $0 \leq \hat{g}_t \leq 1$, the private sector sets $\gamma_t = \hat{g}_t$, for $\hat{g}_t < 0$ it sets $\gamma_t = 0$, and for $\hat{g}_t > 1$ it sets $\gamma_t = 1$.

Interpreted literally, the suggestion that the private sector may practice constant-gain learning because of the possibility of structural change implies that κ^v should equal κ^f. But we allow for the possibility that $\kappa^v \neq \kappa^f$ to explore circumstances in which the ability of the private sector to assess the quality of central bank forecasts may differ from its ability to produce its own forecasts.[10] A small value of κ^f can be seen as a high ability to assess the quality of the central bank's forecast. This is a straightforward interpretation here since there is no structural change taking place.[11] A large value of κ^f, on the other hand, can be seen as a low ability to assess the quality of the central bank's forecasts. This will be associated with more variation in γ_t over time and will reduce the accuracy of the combined forecast.

3. Results

We start by exploring the limiting case in which the forecasting models used by both the central bank and the private sector accurately describe the dynamics of inflation. In this case, the central bank and the private sector arrive independently at identical inflation forecasts

[10] Regarding terminology, the use of constant-gain learning should here primarily be seen as an active response to the limited-information problem faced by the private sector; an alternative interpretation is that it is a matter of finite memory, that is, a deviation from full rationality.

[11] Outside the rather narrow framework of the model, κ^f could also be seen as partly reflecting the extent to which the central bank makes and communicates changes in its analytical framework. For example, replacing or respecifying the main model used for analysis or changing the forecasting methods in other ways—such as introducing a suite of models or taking a different approach to judgmental adjustments—is likely to affect the private sector's opinion about how to weight the forecasts together.

and central bank communication has no role to play. This limiting case, which mimics the rational expectations, perfect-knowledge outcome, serves as a benchmark against which to consider the impact of imperfect knowledge and central bank communications. We next turn to the intermediate case in which the central bank has perfect knowledge about the structure of the economy and uses its forecasts to help inform the private sector. The benefits of central bank communications depend on the ability of agents to assess the quality of the information being conveyed. We then explore the role of central bank communications in our main case, in which both the central bank and the private sector have imperfect knowledge about the structure of the economy. We first explore the risk of distraction in the case in which the central bank has only one piece of information which it is able to communicate—its inflation forecast. We show that publishing the central bank's forecast has an ambiguous impact on the accuracy of private-sector expectations depending on the relative precision of the central bank and private-sector forecasts. We then consider the value of publishing the central bank forecast in the context of a communications strategy in which the central bank is also able to announce its inflation target. We show that announcing an inflation target may "crowd out" a role for publishing the central bank's economic forecasts.

In the baseline simulations, the parameters of the supply function are set to $\phi = \alpha = 0.5$ and we let the error term have variance $\sigma_e^2 = 1$. The responsiveness of monetary policy (in form of the output gap) to the inflation gap is set to $\theta = 0.6$.[12] The gains used by the private sector in equations (7) and (12) are set to $\kappa^v = \kappa^f = 0.03$.[13] The effectiveness of central bank communication is evaluated by

[12] A value of $\theta = 0.6$ would be close to optimal in the case of full information and rational expectations if the central bank was trying to minimize a loss function of the type $L = \omega \text{var}(\pi_t - \pi^*) + (1 - \omega)\text{var}(y_t)$ with $\omega = 0.5$; see, for example, Orphanides and Williams (2004).

[13] Recall that the model is highly stylized and the results are meant only to be illustrative. However, to aid interpretation, the calibration of the supply function ($\phi = \alpha = 0.5$) is similar to estimates reported by Orphanides and Williams (2006) using quarterly U.S. data. To the extent that the model can be interpreted as quarterly, a value, for example, of $\kappa^v = 0.03$ implies that the private sector bases its model estimates on roughly sixteen years of data. See Orphanides and Williams (2004, 2006) for a discussion of the interpretation of, and plausible values for, private-sector gain.

the accuracy of private-sector inflation expectations, measured by their root mean squared error (RMSE).[14] The RMSE of the central bank's forecasts, private sector's forecasts, and private sector's expectations are denoted $RMSE_{\pi^{cb}_{t|t-1}}$, $RMSE_{\pi^{v}_{t|t-1}}$, and $RMSE_{\pi^{e}_{t|t-1}}$, respectively. For each combination of parameters, the economy was simulated $T = 160,000$ time periods; the first 80,000 observations were discarded and the analysis was accordingly based on the second half of each sample.

3.1 Perfect-Knowledge Benchmark

The forecasting models used by the central bank and private sector nest the correct structure of the economy that would prevail under full information and rational expectations. The rational expectations, full-information benchmark can hence be obtained by setting the gains used by the central bank and private sector in estimating their forecasting models to be inversely related to the age of the data, $\kappa^{cb} = \kappa^{v} = 1/t$, so as t increases, κ^{cb} and κ^{v} converge to zero.[15] In this case, the estimation algorithms used by the central bank and private sector collapse to more conventional least-squares learning with infinite memory, and thus the estimates of the two forecasting models converge to their correct values and the perfect-knowledge benchmark solution is obtained. The central bank and private sector produce identical forecasts for inflation and, as such, there is no role for central bank communication. As reported in table 1 in appendix 1, the RMSE of private-sector inflation expectations in this case is governed by the variance of the shocks affecting the economy.[16]

3.2 Intermediate Case: Private-Sector Imperfect Knowledge

Consider now the intermediate case in which the central bank's knowledge of the economy is assumed to converge to the

[14] In line with, for example, Honkapohja and Mitra (2003) and Ehrmann and Fratzscher (2007), we argue that predictability can be seen as a proxy for welfare.

[15] Here, t denotes the distance in time between the observation being weighted and the current observation.

[16] The RMSEs are very close to the true value of σ^2_e. The deviation is due to random error—this has been established by choosing different random-number seeds.

full-information case as data accumulate (that is, $\kappa^{cb} = 1/t$), but the private sector is perpetually learning about the economy.

If the central bank does not publish its forecast, the private sector bases its expectations solely on its own least-squares forecasts from equation (10), such that $\pi^e_{t+1|t} = \pi^v_{t+1|t}$. The results presented in table 2 in appendix 1 show how the RMSE of the private-sector forecast, not surprisingly, is increasing with the gain used in its least-squares algorithm. That is, as the private sector restricts the use it makes of historical data in estimating its forecast model—as κ^v increases—the precision of its forecasts (and hence its expectations) deteriorates. The forecasting performance of the central bank also deteriorates as κ^v increases. This reflects the fact that, given the structure of the model, the forecasting performance of the central bank and private sector will typically depend on each other.

Alternatively, if the central bank publishes its forecasts, the private sector now needs to recursively update estimates of both its forecasting model (7) and the weight to attach to the central bank's forecast when forming its inflation expectations. Table 3 in appendix 1 shows the effect of publishing the central bank's forecast on the RMSE of private-sector inflation expectations when $\kappa^v = \kappa^f$. Comparing the outcomes reported in tables 2 and 3, it can be seen that it is unambiguously beneficial for the central bank to publish its forecast. The central bank has an informational advantage which it can use to help inform private-sector expectations.

The extent of that benefit depends on the private sector's ability to recognize the true value of this forecast. This can be illustrated by varying the gain κ^f used by the private sector to estimate the weight to attach to the central bank's forecast. For simplicity, the gain used by the private sector to estimate its forecasting model (7) is held constant at $\kappa^v = 0.03$. Table 4 in appendix 1 shows that as the private sector's ability to evaluate the quality of the central bank's information improves—that is, the value of κ^f falls—the private sector attaches increasing weight to the central bank's forecast and the precision of its inflation expectations correspondingly improves.[17]

[17]Figure 7 in appendix 1 shows how the weight γ_t evolves over 300 time periods in a representative simulation where $\kappa^v = \kappa^f = 0.03$. As can be seen, γ_t moves around a fair bit and hits both 0 and 1. It is substantially more frequent that $\gamma_t = 0$ than $\gamma_t = 1$, though, signaling that the private sector often places a large weight on the central bank's forecast, which here is of good quality.

3.3 *Main Case: Central Bank and Private-Sector Imperfect Knowledge*

We turn now to our central case in which both the central bank and the private sector are assumed to be perpetually learning about the economy. We start by continuing to consider the case in which the central bank is able to communicate only its inflation forecast. That the central bank also has imperfect knowledge about the economy means that the central bank's forecasts may not necessarily be more precise than those of the private sector. This gives rise to the possibility that if the central bank's forecast is relatively noisy and the private sector inadvertently places too much weight on this forecast, publishing the central bank's forecast may detract from the accuracy of private-sector expectations: It may act as a source of distraction.

The effect of publishing central bank forecasts on the RMSE of private-sector expectations is shown by the upward-sloping line in figure 2.[18] As in the intermediate case, the private sector is assumed to update its forecasting equation (7) each period using constant-gain learning and forms its inflation expectations by combining its own inflation forecast with the central bank's forecast; in doing this, the gains κ^v and κ^f are both fixed at 0.03. The effect of varying the quality of the central bank's forecast is illustrated by varying the gain used by the central bank, κ^{cb}; as κ^{cb} increases, the central bank makes less use of historical data to estimate its forecasting model and so the quality of its forecasts deteriorates.[19] By way of comparison, the RMSE of private-sector expectations in the case in which the central bank does not publish its inflation forecast is shown by the horizontal line.

[18] More detailed results are given in tables 5 and 6 in appendix 1.

[19] It should be noted though that very small values of κ^{cb} in some cases generate an RMSE for the central bank which is lower than that which is associated with $\kappa^{cb} = 0$. This finding—while perhaps not completely intuitive—is similar to that of Orphanides and Williams (2004), who find that using more than one lag in the AR model can be beneficial for forecasting performance even though only one lag of inflation appears in the equations for inflation and inflation expectations. Learning introduces time variation in the formation of inflation expectations and thereby in the processes of inflation and the output gap; this feature of the model means that central bank learning with a low gain can be beneficial from a forecasting point of view (relative to knowing α). This is most clearly illustrated in figure 9 in appendix 2.

Figure 2. Private-Sector RMSEs under Private-Sector Imperfect Knowledge (Central Bank Communicates Nothing or Inflation Forecasts: Effect of Varying the Central Bank Gain)

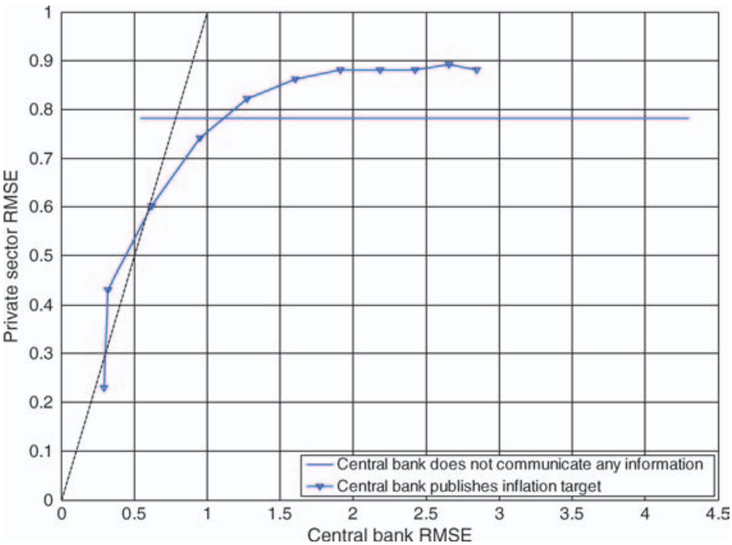

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Private-sector gains are 0.03. Responsiveness of the output gap to the inflation gap is 0.60. RMSEs are given as percentage deviations from the perfect-knowledge benchmark value of 0.998. The dashed straight black line is the 45-degree line. Each triangle represents a value of the central bank gain which has been investigated; these range from 0.00 to 0.10.

Not surprisingly, when the central bank's forecasts are relatively accurate—that is, the central bank gain κ^{cb} is relatively low—publishing the central bank's forecasts improves the accuracy of private-sector expectations relative to the case in which they are not published. However, if the accuracy of the central bank's forecast deteriorates beyond a certain point, then it is better for the central bank not to publish its forecast. That is, the RMSE of private-sector expectations is lower in the case in which the forecast is not published. The possibility that publishing central bank forecasts may detract from the accuracy of private-sector expectations stems from the imperfect ability of the private sector to assess the quality of the forecasts. The private sector is perpetually having to learn about the relative quality of the central bank's forecasts and so there is a risk that it may attach too little or too much weight to the central bank forecast relative to its own forecast. If the private sector attaches too little weight, publishing the central bank forecast still helps to improve the accuracy of private-sector expectations, but not by its full potential. In contrast, if the private sector attaches too much weight to the central bank forecast, there is a risk that it will detract from the accuracy of its expectations.[20]

The point at which publishing the central bank's forecast may act as a source of distraction rests on two key factors. First, it depends on the precision of the central bank's forecasts relative to that of the private sector. As the relative quality of the central bank's forecasts deteriorates, there is a greater chance that publishing the forecasts will distract the private sector.[21] Figure 3 and table 5 (in appendix 1) illustrate the effect of publishing the central bank's forecasts for three different levels of precision of private-sector forecasts (proxied by varying the gain used by the private sector in the algorithm used to estimate its forecasting model: $\kappa^v = (0.01 \quad 0.03 \quad 0.05)$).[22] As the precision of private-sector forecasts improves—that is, the value of κ^v declines—the minimum level of accuracy necessary for it to be beneficial for the central bank to publish its forecast becomes more strenuous.[23]

[20]It should be noted that the risk of distraction is not an "asymptotic" property but can happen also at small sample sizes. A very large number of observations is used in the simulations in order to remove random noise from the results. That is, by simulating the economy for many periods, we get a good estimate of, for example, the true RMSEs for a given parameterization.

[21]As the relative quality of the private-sector forecast increases, the "optimal" weight to place on the central bank forecast falls. Given that the weight that can be placed on the central bank's forecast has a lower bound of zero, this implies that as the relative quality of the private-sector forecast increases, the distribution of γ_t around the "optimal" weight becomes increasingly positively skewed.

[22]For simplicity, the ability of the private sector to assess the relative quality of the central bank and private-sector forecasts (proxied by κ^f) is held constant at $\kappa^f = 0.03$ in all three cases. The case $\kappa^v = \kappa^f = 0.03$ is illustrated graphically in figure 2.

[23]In the case of $\kappa^v = 0.05$, the relatively poor quality of private-sector forecasts means that it is beneficial to publish central bank forecasts even when the central bank gain increases to $\kappa^{cb} = 0.10$.

Figure 3. Private-Sector RMSEs under Private-Sector Imperfect Knowledge (Central Bank Communicates Nothing or Inflation Forecasts: Effect of Varying Gain Used to Estimate AR(1) Model and Central Bank Gain)

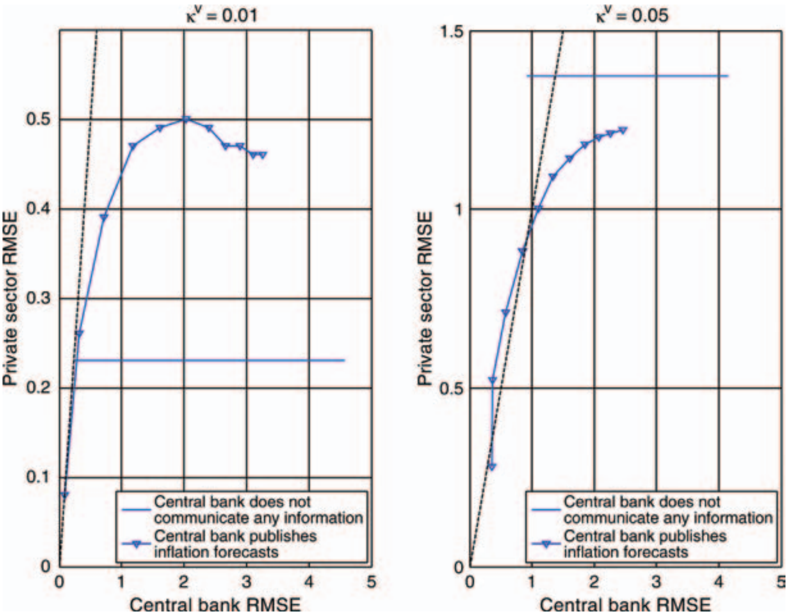

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Private-sector gain used to combine forecasts is 0.03. Responsiveness of the output gap to the inflation gap is 0.60. RMSEs are given as percentage deviations from the perfect-knowledge benchmark value of 0.998. The dashed straight black line is the 45-degree line. Each triangle represents a value of the central bank gain which has been investigated; these range from 0.00 to 0.10.

Second, the value of the central bank publishing its forecasts depends on the ability of the private sector to evaluate their quality. If the private sector's ability to assess the true quality of the imperfect information being communicated is relatively limited, there is a greater chance that publishing noisy economic forecasts will mislead the private sector. This is shown in figure 4 and table 6 (in appendix 1), which illustrate the effect of varying the gain κ^f used by the private sector to estimate the weight to attach to the

Figure 4. Private-Sector RMSEs under Private-Sector Imperfect Knowledge (Central Bank Communicates Nothing or Inflation Forecasts: Effect of Varying Gain Used to Combine Forecasts and Central Bank Gain)

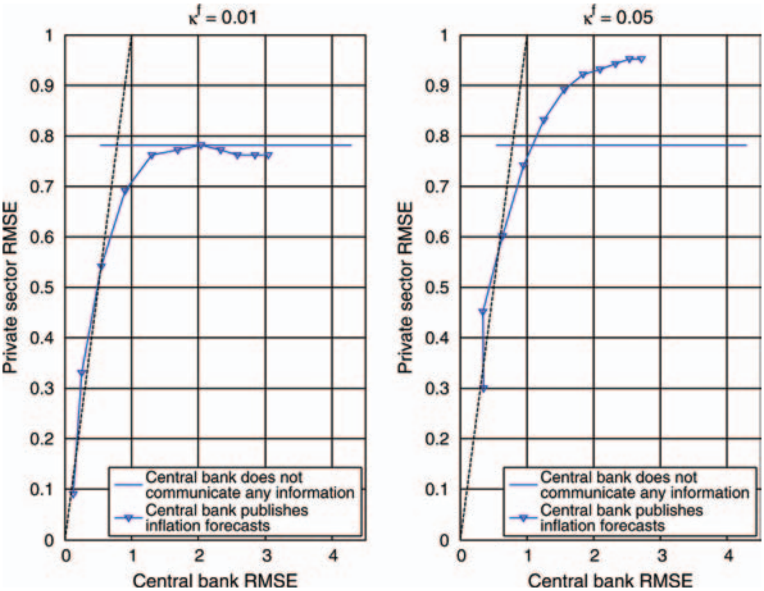

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Private-sector gain used to estimate AR(1) model is 0.03. Responsiveness of the output gap to the inflation gap is 0.60. RMSEs are given as percentage deviations from the perfect-knowledge benchmark value of 0.998. The dashed straight black line is the 45-degree line. Each triangle represents a value of the central bank gain which has been investigated; these range from 0.00 to 0.10.

central bank's forecasts.[24] The left panel of figure 4 shows that the distraction problem is alleviated when the private sector's ability to assess the quality of the central bank's forecast is good, whereas the right panel shows that the problem is exacerbated when the private sector's assessment ability is poor.

[24]Three different values are considered—$\kappa^f = (0.01 \quad 0.03 \quad 0.05)$—while κ^v is kept fixed at 0.03. The case $\kappa^v = \kappa^f = 0.03$ is illustrated graphically in figure 2.

It can be noted that our findings are qualitatively similar to those of Morris and Shin (2002), who show that the welfare effects of communicating public information are ambiguous. The mechanisms at work are different though; in this paper, the risk that communicating imperfect information may act as a source of distraction arises because the private sector is not able to assess correctly the relative quality of that information. This raises the question of whether central banks are able to provide greater guidance about the relative quality of the information they communicate. Indeed, it is notable that central banks in recent years have made considerable efforts to convey the uncertainty associated with some of their communications, such as with the use of so-called “fan charts” to illustrate the subjective probability distribution associated with their forecasts. But although a central bank may be able to make estimates about the *absolute* precision of some of its communications, that is very different from being able to provide estimates of the *relative* precision of that information, which is likely to vary substantially across different sectors of society. And the ability of central banks to form estimates of even the absolute precision of much of the qualitative information they communicate, such as speeches by policymakers or the minutes of policy meetings, is likely to be very limited. Thus, while efforts by central banks to convey the uncertainty of the information they communicate may help to mitigate this problem, it is unlikely to represent a complete solution. Moreover, when considering the relative precision of the marginal piece of information that can be communicated, it is important that it is evaluated in the context of the information that is already being conveyed, that is, in terms of the central bank’s existing communications strategy. We turn to this issue next.

3.4 Communication Strategies

So far we have assumed that the central bank has only one piece of information that it can potentially communicate, namely its inflation forecasts. Thus the central bank’s communication strategy boils down to a binary decision of whether or not to communicate. But in practice all central banks communicate some information. The challenge faced by central banks is to design a communications strategy which combines various types of information in a way that helps to inform private-sector expectations in an efficient and effective manner. The importance of analyzing central bank communications in the context of a communications strategy can be illustrated by considering the case in which the central bank has the ability to announce its inflation target (π^*) as well as publish its inflation forecast.

To the private sector, knowledge of the inflation target means that it no longer has to estimate the intercept in its econometric model.[25,26] It can accordingly estimate the restricted model

$$\pi_t - \pi^* = c_{1,t}(\pi_{t-1} - \pi^*) + \upsilon_t \tag{13}$$

and forecasts based on this constant-gain least-squares estimation are generated as

$$\pi^{\upsilon}_{t+1|t} = (1 - \hat{c}_{1,t})\pi^* + \hat{c}_{1,t}\pi_t. \tag{14}$$

The private sector continues to form its inflation expectations by combining its private forecasts with the central bank's forecasts in the same way as that described earlier.

Figure 5 considers the effect of four alternative communication strategies on the RMSE of private-sector expectations.[27] As before, the private sector is assumed to be perpetually learning about both the structure of the economy and the relative quality of the central

[25]Imposing that the inflation target is the unconditional mean of inflation implicitly rests on the assumption that the public knows the shape of the reaction function and the modified Lucas supply function. This assumption can of course be questioned given the imperfect-knowledge environment introduced. However, we believe that the clarity it brings to the exposition justifies its usage. It is well known, for example, that an asymmetric central bank loss function can generate outcomes that on average are below (or—perhaps more unlikely given the way we think about central bank preferences in general—above) the target. While it could be argued that such aspects have some empirical relevance, where the formulation of the European Central Bank's inflation target might constitute one example, they would add complexity to the model and reduce its transparency without contributing much to our general understanding.

[26]For simplicity, we assume that the announcement of the inflation target is perfectly credible. Allowing for the possibility of imperfect credibility—for example, by letting only a share of the private sector impose the restriction—would not affect the qualitative results.

[27]Results are also shown in tables 7 and 8 in appendix 1.

Figure 5. Private-Sector RMSEs under Central Bank and Private-Sector Imperfect Knowledge for Different Communication Strategies: Effect of Varying the Central Bank Gain

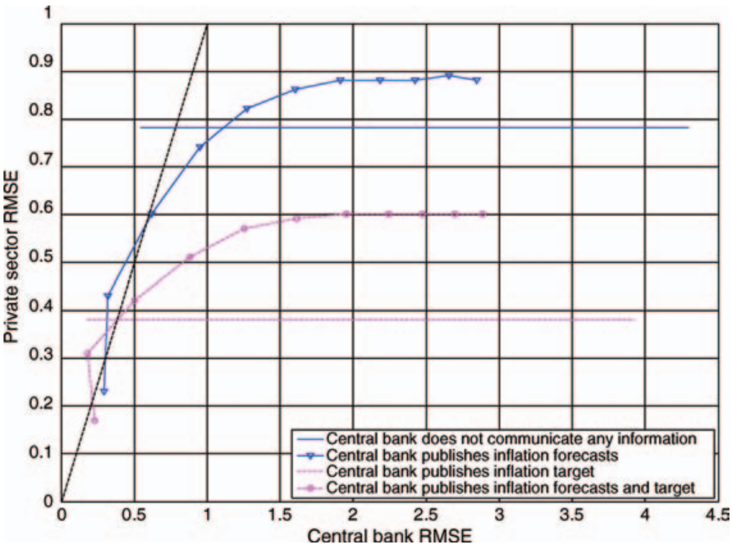

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Private-sector gains are 0.03. Responsiveness of the output gap to the inflation gap is 0.60. RMSEs are given as percentage deviations from the perfect-knowledge benchmark value of 0.998. The dashed straight black line is the 45-degree line. Each triangle/circle represents a value of the central bank gain which has been investigated; these range from 0.00 to 0.10.

bank's forecasts; the gains used by the private sector are set to $\kappa^v = \kappa^f = 0.03$. The upper two lines simply repeat the first results regarding communication strategies considered in section 3.3. The horizontal solid line shows the case in which the central bank does not communicate any information, and the upward-sloping line with triangles shows the case in which the central bank publishes only its inflation forecasts. The lower two lines illustrate the effect of communication strategies which involve announcing an inflation target. The horizontal dashed line shows the case in which the central bank announces only its inflation target, and the upward-sloping line with circles considers the outcome of a communications strategy in which the central bank both announces its inflation target and publishes its inflation forecasts.

Not surprisingly, the communication strategies which include the announcement of an inflation target are unambiguously better than those which do not: Communicating information known with certainty, such as an inflation target, means there is no risk that the private sector will place too much weight on that information.[28] Moreover, comparing the communication strategies which involve announcing an inflation target, it is clear that the benefit of publishing the central bank's inflation forecast remains ambiguous. Once the accuracy of the central bank's forecasts falls below a certain level, it is better for the central bank to announce only its inflation target, rather than also publish its inflation forecast.

The importance of viewing central bank communications in the context of an overall communications strategy can be seen by considering the central bank's decision of whether or not to publish its inflation forecast. In particular, the minimum standard of accuracy required for it to be beneficial for the central bank to publish its inflation forecasts is more strenuous in the case in which the central bank announces an inflation target than in the case in which it does not. This is shown in figure 5 by the fact that the level of central bank gain κ^{cb} at which the forecasts start to detract from the accuracy of private-sector expectations is lower when the central bank announces its inflation target than in the case in which it does not. The intuition for this result stems directly from the fact that the benefits (and costs) of publishing the central bank's forecasts depend on the *relative* accuracy of the central bank and private-sector forecasts. By announcing its inflation target, the central bank improves the accuracy of the private sector's forecasts and so increases the risk that publishing its inflation forecasts will act as a source of distraction. Put more generally, the ability of the central

[28]This assumes that the central bank is genuinely committed to achieving the announced inflation target. It also assumes that the private sector's knowledge of the economy is such that it is able to utilize this information correctly. It is possible that the private sector may place too little weight on the inflation target; that is, the target is not perfectly credible. But, as noted above, this possibility does not affect the qualitative results.

bank to communicate more certain information has the potential to "crowd out" a role for communicating imperfect information.

The results discussed so far have stressed the *qualitative* implications of the model. The highly stylized nature of the model limits the inferences that can usefully be drawn about its *quantitative* implications. In particular, the simplistic nature of the model means that the results should not be interpreted as applying literally to the relative forecasting performance of central banks. The model ignores many of the channels through which central bank economic forecasts may help to guide and inform private-sector expectations. However, it is possible to use the ratio of the RMSE of the central bank and private-sector forecasts as a rough gauge of the relative precision at which it may become potentially harmful for a central bank to communicate imperfect information. To that extent, the results presented in figure 5 suggest that central bank information need only be slightly less precise than that of the private sector for it to be potentially harmful to communicate. For example, the upward-sloping line with circles in figure 5 crosses the horizontal dashed line close to the 45-degree line (which describes equal precision of private-sector and central bank forecasts). Thus, in the context of Svensson's (2006) response to Morris and Shin (2002), the results suggest that the information being communicated by a central bank need not diverge very far from the "conservative benchmark" of equal precision for it to risk distracting the private sector.

Regarding the crowding out, it can be noted that it is a robust finding and is not particularly sensitive to the parameterization of the model. This can be illustrated by varying the responsiveness of the output gap to the inflation gap, that is, the value of θ.[29] The results for $\theta = 0.40$ are summarized in figure 6. The qualitative results are broadly similar to the benchmark case discussed above when $\theta = 0.60$, although it is interesting to note that reducing the responsiveness of the output gap to the inflation gap tends to improve the relative performance of the central bank's forecasts. As a result, for the range of central bank gains considered here, it is

[29] Varying the value of ϕ would affect the dynamics of inflation in a qualitatively similar way.

Figure 6. Private-Sector RMSEs under Central Bank and Private-Sector Imperfect Knowledge for Different Communication Strategies: Effect of Varying the Central Bank Gain (Responsiveness of the Output Gap to the Inflation Gap Is 0.40)

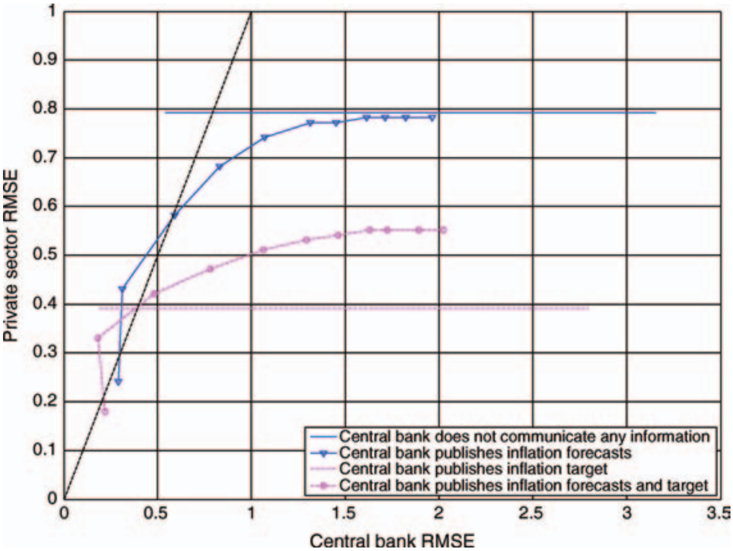

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Private-sector gains are 0.03. RMSEs are given as percentage deviations from the perfect-knowledge benchmark value of 0.998. The dashed straight black line is the 45-degree line. Each triangle/circle represents a value of the central bank gain which has been investigated; these range from 0.00 to 0.10.

always beneficial for the central bank to publish its forecasts relative to an alternative strategy of no communication. If we allow for the possibility that the central bank may be able to announce its inflation target, the benefit of the central bank publishing economic forecasts is once again ambiguous. Setting $\theta = 0.83$, the results are qualitatively similar to those in the benchmark case: Publishing the central bank's forecasts risks distracting the private sector relative to both a no-communication strategy and a strategy of only announcing the inflation target. (Results are not reported but are available upon request.)

An important implication of this crowding-out result is that as more information is communicated by the central bank—and absorbed by the private sector—there is an increasing risk that the marginal piece of information conveyed by the central bank may act as a source of distraction.[30] This will be the case even if the quality of the information being communicated does not deteriorate. If, as seems likely, the precision of the marginal piece of information deteriorates as central banks communicate more information, this further heightens the risk that communicating ever-increasing amounts of information may act as a source of distraction.

4. Conclusions

Mystique and secrecy are truly no longer the bywords of central banking. The importance of open and transparent policy processes is widely recognized and understood. Central banks have made huge strides in their communication policies over the past two decades. But for many central banks, communications remain a "work in progress" (Bernanke 2007). There is a need to better understand the design and evaluation of central bank communication policies. The mantra of "more information is always better" is neither sufficient nor correct.

This paper considers the role of central bank communications in a model of imperfect knowledge and learning. The recognition that monetary policy is conducted in an environment of imperfect information is central to understanding both the potential benefits and limitations of central bank communications. It rationalizes the role central bank communications may play in helping to inform private-sector decisions and expectations. But it also serves to emphasize that much of the information communicated by central banks is

[30]This stems from the recognition that central bank communication can aid private-sector learning; that is, the accuracy of the private sector's forecasts depends on the information communicated by the central bank. In contrast, under the heterogeneous-information assumptions employed by Morris and Shin (2002) among others, the precision of the private sector's information is independent of the information communicated by the central bank. The endogeneity of the private sector's understanding generated by this learning mechanism underpins the importance of considering central bank communication in the context of the overall communications strategy.

noisy and imperfect. Such imperfect information can inform and improve the public's understanding. However, unless interpreted correctly, it also has the potential to distract and mislead; it is shown that the lower the private sector's ability to assess the quality of the central bank's forecast, the likelier it is that noisy information will distract the private sector. Imperfect central bank communication is a double-edged sword that should be used with care.

The possibility that central bank communications may detract from the public's understanding is most apparent when viewed in the context of a communications strategy comprising numerous pieces of information. As more information is conveyed by central banks—and understood by the public—the benefits of increased communication have to be weighed against the growing risk that they will act as a source of distraction.

The central policy message of this paper is that there may be costs—as well as benefits—associated with publishing ever-increasing amounts of uncertain and noisy information. If the information is too noisy relative to the private sector's existing understanding, it may be better for the central bank not to communicate it. This message may strike many policymakers and central bank officials as little more than common sense—we agree. But it is important to understand and demonstrate the mechanisms that may give rise to it. Communicating imperfect information may be detrimental because of the limited ability of the private sector to assess correctly the quality of that information. This suggests that central banks should focus their communication policies on the information they know most about.

The model considered in this paper is highly stylized, and the results should not be interpreted as literally suggesting that publishing central bank forecasts is, as a practical matter, likely to mislead the public. Indeed, the fact that central banks in every advanced economy now publish some form of economic forecasts suggests that experience has led central banks to conclude that publishing information about their forecasts is useful in informing private-sector decisions and expectations. However, the risk of distraction may help to rationalize why many central banks tend to limit the amount of forecast information they publish to two or three key variables rather than provide detailed numerical forecasts for a large number of variables.

Appendix 1. Tables and Figure 7

Table 1. RMSEs under Perfect-Knowledge Benchmark

$RMSE_{\pi^{e}_{t\|t-1}}$	$RMSE_{\pi^{v}_{t\|t-1}}$	$RMSE_{\pi^{cb}_{t\|t-1}}$
0.998	0.998	0.998

Notes: The central bank and private sector both know all parameters of the model. Responsiveness of the output gap to the inflation gap is 0.60.

Table 2. RMSEs under Private-Sector Imperfect Knowledge (Central Bank Does Not Communicate Any Information)

κ^v	$RMSE_{\pi^{e}_{t\|t-1}}$	$RMSE_{\pi^{v}_{t\|t-1}}$	$RMSE_{\pi^{cb}_{t\|t-1}}$
0.01	1.000	1.000	1.000
0.03	1.006	1.006	1.006
0.05	1.012	1.012	1.012

Notes: The central bank knows all parameters of the model. Responsiveness of the output gap to the inflation gap is 0.60.

Table 3. RMSEs under Private-Sector Imperfect Knowledge (Central Bank Publishes Inflation Forecasts: Effect of Varying Private-Sector Gains)

$\kappa^v = \kappa^f$	$\bar{\gamma}_t$	$RMSE_{\pi^{e}_{t\|t-1}}$	$RMSE_{\pi^{v}_{t\|t-1}}$	$RMSE_{\pi^{cb}_{t\|t-1}}$
0.01	0.247	0.999	1.002	0.999
0.03	0.250	1.000	1.010	1.001
0.05	0.249	1.002	1.019	1.003

Notes: The central bank knows all parameters of the model. Responsiveness of the output gap to the inflation gap is 0.60.

Table 4. RMSEs under Private-Sector Imperfect Knowledge (Central Bank Publishes Inflation Forecasts: Effect of Varying Forecast Combination Gain)

κ^f	$\bar{\gamma}_t$	$RMSE_{\pi^e_{t\|t-1}}$	$RMSE_{\pi^v_{t\|t-1}}$	$RMSE_{\pi^{cb}_{t\|t-1}}$
0.01	0.189	0.999	1.011	0.999
0.03	0.250	1.000	1.010	1.001
0.05	0.275	1.001	1.010	1.002

Notes: The central bank knows all parameters of the model. Private-sector gain used to estimate the AR(1) model is 0.03. Responsiveness of the output gap to the inflation gap is 0.60.

Table 5. RMSEs under Central Bank and Private-Sector Imperfect Knowledge (Central Bank Publishes Inflation Forecasts: Effect of Varying Gain Used to Estimate AR(1) Model and Central Bank Gain)

κ^v	κ^{cb}	$\bar{\gamma}_t$	$RMSE_{\pi^e_{t\|t-1}}$	$RMSE_{\pi^v_{t\|t-1}}$	$RMSE_{\pi^{cb}_{t\|t-1}}$
0.01	0.00	0.301	0.999	1.002	0.999
0.01	0.02	0.610	1.002	1.002	1.005
0.01	0.04	0.698	1.003	1.003	1.014
0.01	0.06	0.715	1.003	1.003	1.022
0.01	0.08	0.718	1.003	1.004	1.027
0.01	0.10	0.722	1.003	1.003	1.031
0.03	0.00	0.250	1.000	1.010	1.001
0.03	0.02	0.312	1.004	1.007	1.004
0.03	0.04	0.498	1.006	1.008	1.011
0.03	0.06	0.582	1.007	1.008	1.017
0.03	0.08	0.610	1.007	1.009	1.022
0.03	0.10	0.630	1.007	1.009	1.026
0.05	0.00	0.222	1.001	1.019	1.002
0.05	0.02	0.230	1.005	1.016	1.004
0.05	0.04	0.334	1.008	1.015	1.009
0.05	0.06	0.430	1.009	1.015	1.014
0.05	0.08	0.485	1.010	1.016	1.019
0.05	0.10	0.515	1.010	1.017	1.023

Notes: Private-sector gain used to combine forecasts is 0.03. Responsiveness of the output gap to the inflation gap is 0.60.

Table 6. RMSEs under Central Bank and Private-Sector Imperfect Knowledge (Central Bank Publishes Inflation Forecasts: Effect of Varying Gain Used to Combine Forecasts and Central Bank Gain)

κ^f	κ^{cb}	$\bar{\gamma}_t$	$RMSE_{\pi^e_{t\|t-1}}$	$RMSE_{\pi^v_{t\|t-1}}$	$RMSE_{\pi^{cb}_{t\|t-1}}$
0.01	0.00	0.189	0.999	1.011	0.999
0.01	0.02	0.276	1.003	1.008	1.004
0.01	0.04	0.544	1.006	1.007	1.011
0.01	0.06	0.645	1.006	1.008	1.018
0.01	0.08	0.676	1.006	1.008	1.024
0.01	0.10	0.695	1.006	1.008	1.028
0.03	0.00	0.250	1.000	1.010	1.001
0.03	0.02	0.312	1.004	1.007	1.004
0.03	0.04	0.498	1.006	1.008	1.011
0.03	0.06	0.582	1.007	1.008	1.017
0.03	0.08	0.610	1.007	1.009	1.022
0.03	0.10	0.630	1.007	1.009	1.026
0.05	0.00	0.275	1.001	1.010	1.002
0.05	0.02	0.323	1.004	1.008	1.004
0.05	0.04	0.476	1.006	1.008	1.011
0.05	0.06	0.550	1.007	1.009	1.016
0.05	0.08	0.578	1.007	1.010	1.021
0.05	0.10	0.597	1.008	1.010	1.025

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Private-sector gain used to estimate the AR(1) model is 0.03. Responsiveness of the output gap to the inflation gap is 0.60.

Table 7. RMSEs under Private-Sector Imperfect Knowledge (Central Bank Communicates Inflation Target Only)

$RMSE_{\pi^e_{t\|t-1}}$	$RMSE_{\pi^v_{t\|t-1}}$	$RMSE_{\pi^{cb}_{t\|t-1}}$
1.002	1.002	1.002

Notes: The central bank knows all parameters of the model. Private-sector gain used to estimated the AR(1) model is 0.03. Responsiveness of the output gap to the inflation gap is 0.60.

Table 8. RMSEs under Central Bank and Private-Sector Imperfect Knowledge (Central Bank Publishes Inflation Target and Inflation Forecasts: Effect of Varying the Central Bank Gain)

κ^{cb}	$\bar{\gamma}_t$	$RMSE_{\pi^e_{t\|t-1}}$	$RMSE_{\pi^v_{t\|t-1}}$	$RMSE_{\pi^{cb}_{t\|t-1}}$
0.00	0.254	1.000	1.003	1.000
0.02	0.552	1.002	1.002	1.003
0.04	0.705	1.004	1.003	1.011
0.06	0.720	1.004	1.003	1.018
0.08	0.721	1.004	1.004	1.023
0.10	0.723	1.004	1.004	1.027

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Private-sector gains are 0.03. Responsiveness of the output gap to the inflation gap is 0.60.

Figure 7. Evolution of γ_t in Standard Simulation under Private-Sector Imperfect Knowledge (Central Bank Publishes Inflation Forecasts)

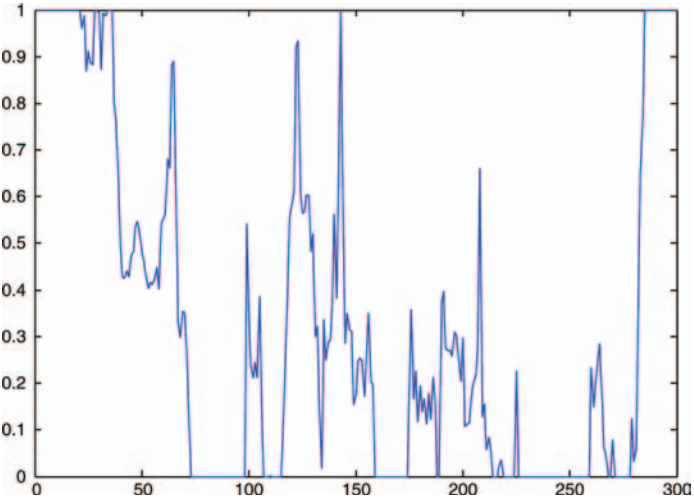

Notes: The central bank knows all parameters of the model. Private-sector gains are 0.03. Responsiveness of the output gap to the inflation gap is 0.60.

Appendix 2. Additional Sensitivity Analysis

As stressed in the main part of the paper, the model considered here is highly stylized, and as such the quantitative results depend on its precise specification and parameterization. To illustrate this sensitivity, we conduct some additional sensitivity analysis in this appendix.

Changing the Private-Sector Gain

Consider the effects of changing the gains used by the private sector. In particular, instead of the benchmark values of $\kappa^v = \kappa^f = 0.03$, consider the effect of the private sector using significantly smaller $\kappa^v = \kappa^f = 0.01$ or larger $\kappa^v = \kappa^f = 0.10$ private-sector gains. The responsiveness of the output gap to the inflation gap is set to the benchmark value of $\theta = 0.6$. Results are reported in figures 8 and 9.

As would be expected, allowing the private sector to reduce its gain to $\kappa^v = \kappa^f = 0.01$ greatly improves the precision of the private sector's inflation expectations; the RMSEs under different communication strategies shown in figure 8 are noticeably lower than those in the corresponding benchmark case. For it to be beneficial for the central bank to publish its forecasts, the central bank's gain now has to be very small.

The reverse is the case when the private sector is forced to use the relatively high gain of $\kappa^v = \kappa^f = 0.10$, shown in figure 9. Indeed, the extent of the private sector's uncertainty means that for the range of central bank gains considered here, it is now always beneficial for the central bank to publish its forecasts.

Changing the Arguments of the Reaction Function

The central bank has in the previous analysis been reacting to inflation outcomes. However, as pointed out by, for example, Evans and Honkapohja (2003), there could be benefits from making the policy rule explicitly depend on private-sector expectations in cases where there are deviations from full information and rational expectations. In line with this reasoning, the central bank's reaction function is next modified so that the private sector's expectations also enter it according to

$$x_t = -\theta\big[\phi\big(\pi^e_{t|t-1} - \pi^*\big) + (1-\phi)(\pi_{t-1} - \pi^*)\big]. \tag{15}$$

Figure 8. Private-Sector RMSEs under Central Bank and Private-Sector Imperfect Knowledge for Different Communication Strategies: Effect of Varying the Central Bank Gain (Private-Sector Gains Are 0.01)

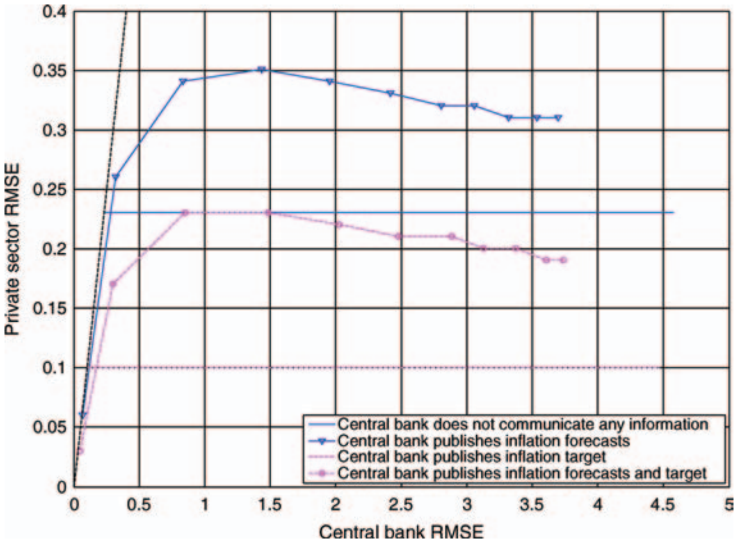

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Responsiveness of the output gap to the inflation gap is 0.60. RMSEs are given as percentage deviations from the perfect-knowledge benchmark value of 0.998. The dashed straight black line is the 45-degree line. Each triangle/circle represents a value of the central bank gain which has been investigated; these range from 0.00 to 0.10.

Results from this exercise are reported in figure 10. As can be seen, the results are qualitatively the same as before. Publishing the inflation target is always beneficial. Publishing the inflation forecasts, however, has a very high probability of distracting the private sector. The central bank's gain once again has to be extremely small for that strategy to reduce the private-sector RMSE.

Letting the Responsiveness of the Output Gap to the Inflation Gap Be Time Varying

The central bank's reaction function in equation (2)—in which the responsiveness of the output gap to the inflation gap is constant—is

Figure 9. Private-Sector RMSEs under Central Bank and Private-Sector Imperfect Knowledge for Different Communication Strategies: Effect of Varying the Central Bank Gain (Private-Sector Gains Are 0.10)

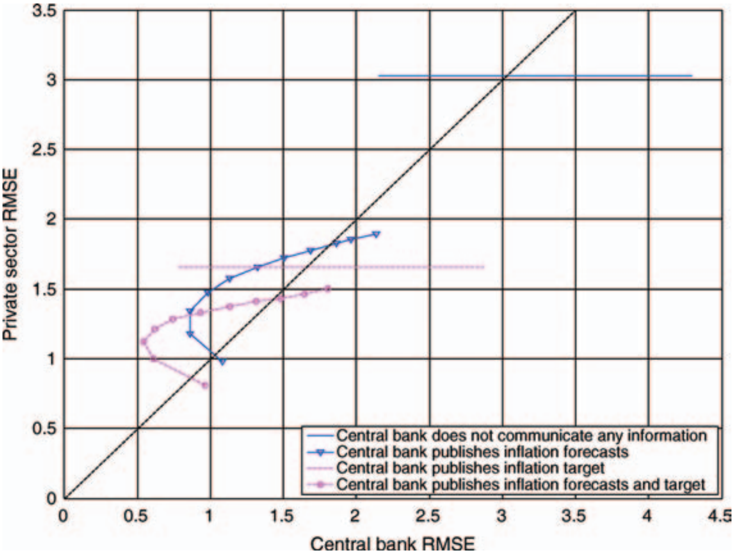

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Responsiveness of the output gap to the inflation gap is 0.60. RMSEs are given as percentage deviations from the perfect-knowledge benchmark value of 0.998. The dashed straight black line is the 45-degree line. Each triangle/circle represents a value of the central bank gain which has been investigated; these range from 0.00 to 0.10.

seen as the central bank employing a simple rule for making policy. This choice is easily motivated, as a simple rule can have several benefits, including robustness.[31] However, in a world with full information and rational expectations, an adequately parameterized version of this rule would in fact be optimal. Following Orphanides and Williams (2004), it can be shown that the optimal value of θ is given by

$$\theta^P = \frac{\omega}{2(1-\omega)} \left(-\frac{\alpha}{1-\phi} + \sqrt{\left(\frac{\alpha}{1-\phi}\right)^2 + \frac{4(1-\omega)}{\omega}} \right), \qquad (16)$$

[31]See, for example, Taylor (1999) and Leitemo and Söderström (2005).

Figure 10. Private-Sector RMSEs under Central Bank and Private-Sector Imperfect Knowledge for Different Communication Strategies: Effect of Varying the Central Bank Gain (Central Bank Reaction Function Has Been Modified to Include Both Actual Inflation and Private Sector's Inflation Expectations)

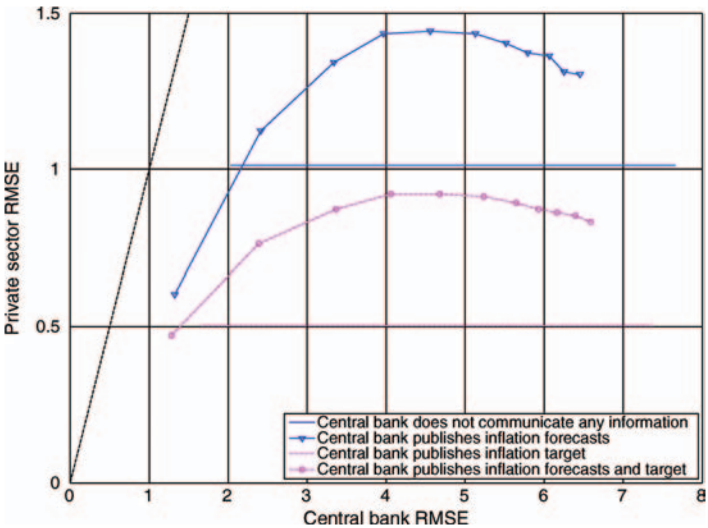

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Private-sector gains are 0.03. Responsiveness of the output gap to the inflation gap is 0.60. Central bank reaction function is given by equation (15). RMSEs are given as percentage deviations from the perfect-knowledge benchmark value of 0.998. The dashed straight black line is the 45-degree line. Each triangle/circle represents a value of the central bank gain which has been investigated; these range from 0.00 to 0.10.

for $0 < \omega < 1$, where ω is the weight on the variance of inflation deviations in the loss function (see footnote 12). In many cases above, we assume that the central bank has to estimate α. But since θ in equation (2) is a function of α, it seems that the central bank might be tempted to let θ be time varying, even though there is little reason to expect this to be a more successful strategy in this setting.

Letting the central bank use equation (16), with α replaced by its estimate $\hat{\alpha}_t$, the central bank's responsiveness becomes time varying. The model is simulated using this setting and the main results

Figure 11. Private-Sector RMSEs under Central Bank and Private-Sector Imperfect Knowledge for Different Communication Strategies: Effect of Varying the Central Bank Gain (Central Bank Lets the Responsiveness of the Output Gap to the Inflation Gap Be Time Varying)

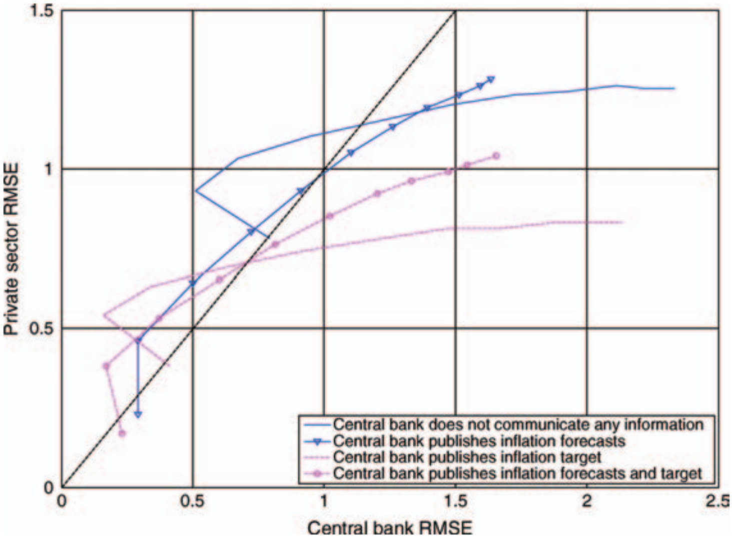

Notes: The central bank estimates elasticity of inflation with respect to the output gap. Private-sector gains are 0.03. Responsiveness of the output gap to the inflation gap in equation (2) is given by equation (16), where the elasticity of inflation with respect to the output gap is replaced by its estimate. RMSEs are given as percentage deviations from the perfect-knowledge benchmark value of 0.998. The dashed straight black line is the 45-degree line. Each triangle/circle represents a value of the central bank gain which has been investigated; these range from 0.00 to 0.10.

are given in figure 11. The main difference from previous results in the paper is that the precision of the private-sector forecasts now changes with κ^{cb} also when nothing, or only the target, is communicated. The reason for this is of course that the central bank affects the behavior of the economy also when it communicates nothing or only the target since the responsiveness of the output gap with respect to the inflation gap is time varying. However, the main conclusions are unchanged: First, it is always beneficial to publish the inflation target. Second, the central bank can distract the private sector by publishing its forecasts. The central bank's gain can be fairly high, though, before publishing inflation forecasts deteriorates private-sector expectations.

References

Bernanke, B. S. 2007. "Inflation Expectations and Inflation Forecasting." Remarks by Ben S. Bernanke, Chairman of the Board of Governors of the Federal Reserve System, at the Monetary Economics Workshop of the NBER Summer Institute, Cambridge, Massachusetts, July 10.

Diebold, F. X., and P. Pauly. 1987. "Structural Change and the Combination of Forecasts." *Journal of Forecasting* 6 (1): 21–40.

Ehrmann, M., and M. Fratzscher. 2007. "Social Value of Public Information: Testing the Limits to Transparency." ECB Working Paper No. 821.

Evans, G. W., and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics.* Princeton, NJ: Princeton University Press.

———. 2003. "Adaptive Learning and Monetary Policy Design." *Journal of Money, Credit, and Banking* 35 (6, Part 2): 1045–72.

Granger, C. W. J., and R. Ramanathan. 1984. "Improved Methods of Combining Forecasts." *Journal of Forecasting* 3 (2): 197–204.

Harrison, R., K. Nikolov, M. Quinn, G. Ramsay, A. Scott, and R. Thomas. 2005. *The Bank of England Quarterly Model.* London: Bank of England.

Honkapohja, S., and K. Mitra. 2003. "Learning with Bounded Memory in Stochastic Models." *Journal of Economic Dynamics and Control* 27 (8): 1437–57.

Issing, O. 1999. "The Eurosystem: Transparent and Accountable or 'Willem in Euroland'?" *Journal of Common Market Studies* 37 (3): 503–19.

———. 2000. "Communication Challenges for the ECB." Speech given at the CFS Conference "The ECB and Its Watchers," Frankfurt, June 26.

King, M. 2006. Speech at the Lord Mayor's Banquet for Bankers and Merchants, London, June 21.

Leitemo, K., and U. Söderström. 2005. "Simple Monetary Policy Rules and Exchange Rate Uncertainty." *Journal of International Money and Finance* 24 (3): 481–507.

Macklem, T. 2005. "Commentary: Central Bank Communications and Policy Effectiveness." Commentary at the Federal Reserve Bank of Kansas City Symposium on The Greenspan Era: Lessons for the Future, Jackson Hole, Wyoming, August 25–27.

Mishkin, F. S. 2004. "Can Central Bank Transparency Go Too Far?" NBER Working Paper No. 10829.

Morris, S., and H. S. Shin. 2002. "Social Value of Public Information." *American Economic Review* 92 (5): 1521–34.

Orphanides, A., and J. C. Williams. 2004. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." In *The Inflation-Targeting Debate*, ed. B. Bernanke and M. Woodford. Chicago: University of Chicago Press.

———. 2006. "Inflation Targeting under Imperfect Knowledge." Working Paper No. 2006-14, Federal Reserve Bank of San Francisco.

Romer, C. D., and D. H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review* 90 (3): 429–57.

Sargent, T. J. 1993. *Bounded Rationality in Macroeconomics.* Oxford: Oxford University Press.

———. 1999. *The Conquest of American Inflation.* Princeton, NJ: Princeton University Press.

Sibert, A. 2009. "Is Transparency about Central Bank Plans Desirable?" *Journal of the European Economic Association* 7 (4): 831–57.

Svensson, L. E. O. 2006. "Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro-Transparency, Not Con." *American Economic Review* 96 (1): 448–52.

Taylor, J. B. 1999. "The Robustness and Efficiency of Monetary Policy Rules as Guidelines for Interest Rate Setting by the European Central Bank." *Journal of Monetary Economics* 43 (3): 655–79.

Winkler, B. 2000. "Which Kind of Transparency? On the Need for Clarity in Monetary Policy-Making." ECB Working Paper No. 26.

Woodford, M. 2005. "Central Bank Communication and Effectiveness." Paper presented at the Federal Reserve Bank of Kansas City Symposium on The Greenspan Era: Lessons for the Future, Jackson Hole, Wyoming, August 25–27.

Inflation Conservatism and Monetary-Fiscal Policy Interactions*

Andrew P. Blake[a] and Tatiana Kirsanova[b]
[a]Bank of England
[b]University of Exeter

This paper investigates the stabilization bias that arises in a model of monetary and fiscal policy stabilization of the economy, when monetary authority puts higher weight on inflation stabilization than society. We demonstrate that inflation conservatism unambiguously leads to social welfare losses if the fiscal authority acts strategically under discretion. Although the precise form of monetary-fiscal interactions depends on the leadership structure, the choice of fiscal instrument, and the level of steady-state debt, the assessment of gains is robust to these assumptions. We develop an algorithm that computes leadership equilibria in a general framework of LQ RE models with strategic agents.

JEL Codes: E31, E52, E58, E61, C61.

1. Introduction

There is a well-understood policy proposal that the agency charged with determining monetary policy, usually the central bank, should be "inflation conservative," by which it is meant that it should have a higher relative weight on the inflation stabilization objective than is socially optimal. The logic runs as follows. Suppose the policymaker

 This paper was completed when Tatiana Kirsanova was a Houblon-Norman/George Research Fellow at the Bank of England. It represents the views and analysis of the authors and should not be thought to represent those of the Bank of England or Monetary Policy Committee members. We are grateful to an editor and an anonymous referee for helpful comments and suggestions. Any errors remain ours. Author addresses: Blake: CCBS, Bank of England, Threadneedle Street, London EC2R 8AH; e-mail: andrew.blake@bankofengland.co.uk. Kirsanova: School of Business and Economics, University of Exeter, Streatham Court, Rennes Drive, Exeter EX4 4PU; e-mail: t.kirsanova@exeter.ac.uk.

must act under discretion. Also suppose, for some reason (and we nearly always assume some political economy issues here), it wants to have lower unemployment (or equivalently higher output). Such a policymaker will generate an excessively high equilibrium *rate* of inflation, a *level bias*. This is the now-textbook Barro and Gordon (1983) model. One possible resolution of this problem is to make the central bank both independent and inflation conservative, as famously shown by Rogoff (1985). This removes some of the incentive to reduce unemployment, and hence reduces the inflation bias. This analysis, and its myriad extensions, has been very influential: most central banks in developed countries are now operationally (or at least target) independent. Even so, central banks are often bestowed with the soubriquet "conservative" in the popular press, even if mitigating the inflationary bias is not quite what journalists have in mind.

There seems little doubt that major central banks do not pursue excessive output targets and none of them aims to stabilize inflation at too high a level. Contemporary policymakers demonstrate considerable agreement about what the *targets* of policy should be. Nonetheless, in *dynamic* models a *stabilization bias* might still arise. If a policymaker acts under discretion, then the *volatility* of welfare-relevant economic variables is necessarily higher than if it acts under commitment; see Currie and Levine (1993). Svensson (1997) calls the difference in the welfare "stabilization bias" and demonstrates, once again, that delegating policy to a conservative central bank could still lead to an improved outcome. Clarida, Galí, and Gertler (1999) show that an inflation-conservative monetary authority that acts under discretion can achieve the same level of welfare as under the optimal precommitment-to-rules policy. In other words, the stabilization bias can be reduced: making the central bank inflation conservative helps solve the problem of optimal delegation.[1] It seems reasonable to conclude that the conservative central bank proposal is a good one, as it deals with both the level and the stabilization bias.

However, it would be naive to assume that an independent inflation-conservative central bank *eliminates* any influence of

[1]The optimal delegation is usually described as a possibility to distort targets of a discretionary policymaker so that it can achieve higher social welfare than it would achieve if it were benevolent.

government, which, after all, still controls the fiscal instruments of financial management. It may deviate from what is socially optimal and, say, have a lower unemployment target or prefer more stable growth. Even if it remains entirely benevolent, it may still behave *strategically*,[2] which can affect the resulting equilibrium. Both level and stabilization bias can still arise.

Indeed, Dixit and Lambertini (2003), using a static Barro-Gordon type model, demonstrate that a conflict of interests does make the outcome suboptimal in a model of monetary-fiscal interactions with a conservative central bank and a *benevolent* fiscal authority. The level bias still arises: the levels of inflation and unemployment in steady-state equilibrium are both higher than socially optimal. With an additional policymaker, it seems the conservative central bank proposal does not work out even in a *static* model.[3] Similarly, the conservative central bank solution to an optimal delegation problem in a *dynamic* stochastic environment could also be misleading if potentially strategic play by policymakers is ignored.[4]

In further related research, Adam and Billi (2008) examine the advantages of inflation conservatism using a non-linear dynamic model similar to ours and find that it can be beneficial. However, they mostly study the implications for the steady state, and limit their analysis of stabilization bias to the case of strategic fiscal leadership and a flexible-price version of the model. By contrast, our paper is concerned with the stabilization aspects of the problem and can be seen as a complement to their paper.

[2]There is little doubt that the fiscal authorities can act strategically: An existing empirical literature on monetary-fiscal interactions suggests that fiscal policy does more than simply allow automatic stabilizers to operate; see, e.g., Auerbach (2003) and Favero and Monacelli (2005), who analyze fiscal policy in the United States. Moreover, since the onset of European monetary union, there are calls for greater fiscal flexibility, although how strategic such authorities should be is not explicitly discussed.

[3]These results were further developed in Lambertini (2006) specifically for the conservative central bank proposal. See also Hughes Hallett, Libich, and Stehlik (2009).

[4]Of course, most policy authorities (the central bank and the fiscal authority) would quite rightly deny that they “play games.” However, by “playing a game” we mean that we model the ability of each authority to understand the other’s priorities and reaction functions.

Making the model dynamic furthers the analysis on two accounts. First, we argue that there is now a wide consensus about appropriate level targets for the financial authorities, so the focus of the policy debate is often *how quickly* should they try to achieve these targets, rather than *which targets* to achieve. Second, modeling monetary-fiscal interactions in a dynamic setting allows us to study the effect of the government's solvency constraint. Debt stabilization issues can impose severe restrictions on the stabilization abilities of both authorities (see Leeper 1991), but it is very difficult to model these restrictions adequately in a static model.[5]

In this paper we explore the importance of both strategic behavior and inflation conservatism in a linear-quadratic rational expectations (LQ RE) infinite-horizon dynamic model of a kind widely used in policy analysis. We employ a conventional model with monopolistic competition and sticky prices in the goods markets (as in Woodford 2003a, ch. 6), extended to include fiscal policy and nominal government debt. We allow both authorities to act strategically and non-cooperatively in pursuit of their own objectives. We then delegate monetary policy to an inflation-conservative central bank. We demonstrate that even if the fiscal authority remains benevolent but acts strategically, delegating monetary policy in this way leads to welfare losses. The basic intuition is that if the authorities' objectives do not coincide, then one strategic policymaker can offset the policy of the other. We show that although a small degree of conservatism can be harmless, greater conservatism leads to substantial welfare losses.

As in Dixit and Lambertini (2003), the quantitative outcome of the game depends on the leadership structure. We study three possibilities: either the monetary or fiscal authorities lead or they move simultaneously (play a Nash game). This analysis is impossible to conduct without developing an appropriate modeling framework. Because an even moderately complicated dynamic model needs to be solved using numerical methods and these methods are neither readily available nor well articulated for models with rational

[5]Beetsma and Bovenberg (2006) introduce public debt into a two-period model. Their results are not readily comparable with ours, as they focus their research on the dynamic of public debt and use an ad hoc welfare metric.

expectations,[6] we develop relevant solution methods for discretionary equilibria in dynamic linear rational expectations models where we make the role of leadership explicit. Although this is a key contribution of the paper and a necessary step in our analysis, we relegate detailed discussion to appendix 1.

This paper differs from Dixit and Lambertini (2003) and Adam and Billi (2008) in four important respects. First, we study the stabilization rather than the level bias, and focus on dynamics rather than the steady state. Second, we examine different leadership regimes, and provide suitable algorithms for calculating appropriate equilibria. In so doing we solve a rather general LQ RE model within a class of dynamic models commonly used by policymakers. Third, we explicitly account for the effects of potential fiscal insolvency. Finally, we show how the interaction of the two authorities depends on the steady-state *level* of debt.

Our results are in marked contrast to the received wisdom outlined above. We find that if the fiscal authority is benevolent but acts strategically, then delegating monetary policy to an inflation-conservative agency usually *increases* stabilization bias and so reduces social welfare. Any distortion to the social objectives can bring the two policymakers into conflict with each other in a way that nearly always reduces social welfare for our model. The message is clear: What works well in an economy with a single policymaker may not work at all in an economy with two strategic policymakers. Our assessment of the gains from delegation seems robust to the specification of the model and the choice of fiscal policy instrument. We conduct sensitivity analysis over the model, and show that whilst there are differences in the behavior of fiscal policy if we choose either distortionary taxes or spending, the qualitative results are the same. However, there is an issue about the steady-state level of debt. The low- and high-debt cases (defined below) have quite different qualitative effects.

The paper is organized as follows. In the next section we present the model and define the private-sector equilibrium. In section 3 we discuss policy design: the choice of instruments, policy objectives,

[6] de Zeeuw and van der Ploeg (1991) provide an excellent discussion of discrete dynamic games which can be compared with our analysis.

the degree of precommitment, and the leadership structure. Calibration of the model is explained in section 4. Section 5 presents the analysis of the benchmark regime with non-cooperative benevolent policymakers, where we compare the outcome of benevolent discretionary policymaking to a solution under commitment with full cooperation of the authorities. In section 6 we study the effect of an inflation conservatism of the monetary authorities on welfare, and we contrast this case to the case of cooperation of benevolent policymakers under discretion. The analysis in these sections assumes the fiscal authorities use spending as a fiscal instrument. Section 7 investigates how our results change if fiscal policy uses tax as an instrument instead or together with spending. Section 8 concludes.

2. The Model

We consider the now-mainstream macro policy model, discussed in Woodford (2003a), modified to take account of the effects of fiscal policy. It is a closed-economy model with two policymakers—the fiscal and monetary authorities. Fiscal policy is allowed to support monetary policy in stabilization of the economy around the steady state.

2.1 Private Sector

2.1.1 Consumers

The economy is inhabited by a large number of households who specialize in the production of a differentiated good (indexed by z) and who spend $h(z)$ of effort in its production. They consume a basket of goods C, and derive utility from per capita government consumption G. The household's maximization problem is

$$\max_{\{C_s,h_s\}_{s=t}^{\infty}} \mathcal{E}_t \sum_{s=t}^{\infty} \beta^{s-t}[u(C_s) + f(G_s) - v(h_s(z))]. \tag{1}$$

The price of a differentiated good z is denoted by $p(z)$, and the aggregate price level is P. All households consume the same basket of goods. Goods are aggregated into a Dixit and Stiglitz (1977) consumption index with the elasticity of substitution between any pair of goods given by $\epsilon_t > 1$ (which is a stochastic[7] elasticity with mean ϵ), $C_t = \left[\int_0^1 c_t^{\frac{\epsilon_t - 1}{\epsilon_t}}(z)dz\right]^{\frac{\epsilon_t}{\epsilon_t - 1}}$.

A household chooses consumption and work effort to maximize criterion (1) subject to the demand system and the intertemporal budget constraint:

$$\mathcal{E}_t \sum_{s=t}^{\infty} Q_{t,s} P_s C_s \leq \mathcal{A}_t + \mathcal{E}_t \sum_{s=t}^{\infty} Q_{t,s} \{(1 - \Upsilon_s)(W_s(z)h_s(z) + \Pi_s(z))) + TR\},$$

where $P_t C_t = \int_0^1 p(z)c(z)dz$ is nominal consumption, $\mathcal{A}_t$ are nominal financial assets of a household, Π_t is profit and TR is a *constant* lump-sum tax or subsidy. Here W is the wage rate, and Υ_t is a tax rate on income. $Q_{t,t+1}$ is the stochastic discount factor which determines the price in period t to the individual of being able to carry a state-contingent amount $\mathcal{A}_{t+1}$ of wealth into period $t+1$. The riskless short-term nominal interest rate i_t is represented in terms of the stochastic discount factor as $\mathcal{E}_t(Q_{t,t+1}) = (1 + i_t)^{-1}$.

We assume that the net present value of future income is bounded and that the nominal interest rate is positive at all times. Optimization for the household requires it to exhaust its intertemporal budget constraint, with wealth accumulation satisfying the no-Ponzi-game condition $\lim_{s \to \infty} \mathcal{E}_t(Q_{t,s}\mathcal{A}_s) = 0$. We assume the specific functional form for the utility-from-consumption component, $u(C_s) = \frac{C_s^{1-1/\sigma}}{1-1/\sigma}$, so household optimization leads to the following dynamic relationship for aggregate consumption:

$$C_t = \mathcal{E}_t \left(\left(\frac{1}{\beta} \frac{P_{t+1}}{P_t} Q_{t,t+1} \right)^{\sigma} C_{t+1} \right). \qquad (2)$$

Additionally, aggregate (nominal) asset accumulation is given by

$$\mathcal{A}_{t+1} = (1 + i_t)(\mathcal{A}_t + (1 - \Upsilon_t)(W_t N_t + \Pi_t) - P_t C_t - TR),$$

where W_t and N_t are aggregate wage and employment.

We linearize equation (2) around the steady state (here and everywhere below for each variable X_t with steady-state value X,

[7] We make this parameter stochastic to allow us to generate shocks to the markup of firms, as did, e.g., Beetsma and Jensen (2004).

we use the notation $\hat{X}_t = \ln(X_t/X)$). Equation (2) leads to the following Euler equation (intertemporal IS curve):

$$\hat{C}_t = \mathcal{E}_t\hat{C}_{t+1} - \sigma(\hat{\imath}_t - \mathcal{E}_t\hat{\pi}_{t+1}). \tag{3}$$

Inflation is $\pi_t = \frac{P_t}{P_{t-1}} - 1$, and we ensure steady-state inflation is zero by appropriate transfers (Woodford 2003a).

2.1.2 Firms

Price setting is based on Calvo contracting as set out in Woodford (2003a). Each period, firms are allowed to recalculate their prices with probability $1 - \gamma$, so that they remain fixed with probability γ. Following Woodford (2003a) we can derive the following Phillips curve for our economy[8]:

$$\hat{\pi}_t = \beta\mathcal{E}_t\hat{\pi}_{t+1} + \frac{(1-\gamma\beta)(1-\gamma)\psi}{\gamma(\psi+\epsilon)}\hat{s}_t, \tag{4}$$

where marginal cost is

$$\hat{s}_t = \frac{1}{\psi}\hat{Y}_t + \frac{1}{\sigma}\hat{C}_t + \frac{\tau}{(1-\tau)}\hat{\Upsilon}_t + \hat{\eta}_t.$$

The shock $\hat{\eta}_t$ is a markup shock and parameter $\psi = v_y/v_{yy}y$. Parameter τ is the steady-state level of tax rate.

Under flexible prices and in the steady state, the real wage is always equal to the monopolistic markup $\mu_t = -(1-\epsilon_t)/\epsilon_t$. Optimization by consumers then implies

$$\frac{\mu^w}{\mu_t} = \frac{y_t^n(z)^{1/\psi}}{\left(1-\hat{\Upsilon}_t^n\right)\left(C_t^n\right)^{1-1/\sigma}}, \tag{5}$$

where superscript n denotes natural levels (see Woodford 2003a), and μ^w is a steady-state employment subsidy which we discuss below. Linearization of (5) and aggregation yields

$$\hat{Y}_t^n\frac{1}{\psi} + \hat{C}_t^n\frac{1}{\sigma} + \frac{\tau}{(1-\tau)}\hat{\Upsilon}_t^n = 0.$$

[8]The derivation is identical to the one in Woodford (2003a), amended by the introduction of markup shocks as in Beetsma and Jensen (2004).

2.2 Government

The government buys consumption goods G_t, taxes income with tax rate Υ_t, raises lump-sum taxes T, pays an employment subsidy μ^w, and issues nominal debt $\mathcal{B}_t$. Debt is assumed to be one period with a risk-free rate of return. The evolution of the nominal debt stock $\mathcal{B}_t$ is

$$\mathcal{B}_{t+1} = (1 + i_t)(\mathcal{B}_t + P_t G_t - \Upsilon_t P_t Y_t - T + \mu^w). \tag{6}$$

The lump-sum taxes (T) and employment subsidy ($\mu^w = T$) are constant and cannot be used to balance the budget or stabilize the economy.

Equation (6) can be linearized as follows:

$$\tilde{b}_{t+1} = \chi \hat{\imath}_t + \frac{1}{\beta}(\tilde{b}_t - \chi\hat{\pi}_t + (1-\theta)\hat{G}_t - \tau(\hat{\Upsilon}_t + \hat{Y}_t)), \tag{7}$$

where we defined a measure of real debt $B_t = \mathcal{B}_t / P_{t-1}$ and denoted the steady-state ratio of debt to output as χ. The steady-state level of debt is determined from the steady-state version of (6), $\chi = \frac{\theta-(1-\tau)}{(1-\beta)}$, and is a function of the steady-state tax rate, τ, and of the steady-state ratio of consumption to output, θ. We denote $\tilde{b}_t = \chi\hat{B}_t$, which allows us to use the same system if $\chi = 0$; then $\tilde{b}_t = B_t$, and there are no first-order effects of either the interest rate or inflation on debt.

2.3 Market Clearing Conditions

Output is distributed as wages and profits:

$$P_t Y_t = W_t N_t + \Pi_t.$$

As government expenditures constitute part of demand, the national income identity is

$$Y_t = G_t + C_t \tag{8}$$

and in steady state $G = (1-\theta)Y$. The linearized national income identity (or the resource constraint) is then

$$\hat{Y}_t = (1-\theta)\hat{G}_t + \theta\hat{C}_t. \tag{9}$$

2.4 Private-Sector Equilibrium

We simplify notation in equations (3), (4), (7), and (9) by using lowercase letters to denote "gap" variables, where the gap is the difference between actual levels and natural levels, i.e., $x_t = \hat{X}_t - \hat{X}_t^n$. The final system of first-order conditions consists of an intertemporal IS curve (10), the Phillips curve (11), national income identity (12), and an equation explaining the evolution of debt (13):

$$c_t = \mathcal{E}_t c_{t+1} - \sigma(i_t - \mathcal{E}_t \pi_{t+1}), \tag{10}$$

$$\pi_t = \beta \mathcal{E}_t \pi_{t+1} + \frac{\kappa}{\sigma} c_t + \frac{\kappa}{\psi} y_t + \frac{\kappa \tau}{(1-\tau)} \tau_t + \eta_t, \tag{11}$$

$$y_t = (1-\theta) g_t + \theta c_t, \tag{12}$$

$$\tilde{b}_{t+1} = \chi i_t + \frac{1}{\beta}(\tilde{b}_t - \chi \pi_t + (1-\theta) g_t - \tau(\tau_t + y_t)), \tag{13}$$

where parameter $\kappa = \frac{(1-\gamma\beta)(1-\gamma)\psi}{\gamma(\psi+\epsilon)}$.

A private-sector rational expectations equilibrium consists of plan $\{c_t, \pi_t, \tilde{b}_t, y_t\}$ satisfying equations (10)–(13), given the policies $\{i_t, g_t, \tau_t\}$, the exogenous process $\{\eta_t\}$, and initial conditions $\tilde{b}_0$.

3. Monetary and Fiscal Policy Regimes

3.1 Welfare Criterion

We assume that both authorities set their instruments to maximize the aggregate utility function:

$$\frac{1}{2} \mathcal{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \left[\frac{C_s^{1-1/\sigma}}{1-1/\sigma} + \xi \frac{G_s^{1-1/\sigma}}{1-1/\sigma} - \int_0^1 \frac{h_s^{1+1/\psi}(z)}{1+1/\psi} dz \right]. \tag{14}$$

We show in appendix 2 that problem (14) implies the following optimization problem for a benevolent policymaker. A benevolent policymaker minimizes the discounted sum of all future losses:

$$L_t = \frac{1}{2} \mathcal{E}_t \sum_{s=t}^{\infty} \beta^{s-t} W_s^S,$$

where the one-period term is

$$W_s^S = \rho_c c_s^2 + \rho_g g_s^2 + \rho_y y_s^2 + \pi_s^2 + \mathcal{O}(\|\xi\|^3) \tag{15}$$

and $\mathcal{O}(\|\xi\|^3)$ collects terms of higher order and terms independent of policy. We normalize the coefficient on inflation to one. This quadratic approximation to social welfare is obtained assuming that there is a steady-state production subsidy $\mu^w = T$ that eliminates the distortion caused by monopolistic competition and income taxes.[9]

Note that expression (15) contains a quadratic term in government spending, g. This term enters the welfare expression because it is assumed in (1) that households derive utility from the consumption of public goods, and that the steady-state level of government spending reflects this. However, if we instead assumed that government spending was pure waste but the government still used g as a policy instrument, then changes in g would still influence social welfare through the national income identity, but it would not constitute an independent source of welfare loss.

Each policymaker minimizes its own loss functions. If it is benevolent, it adopts the social loss function. Note that when assigning social welfare (15) to the monetary authority, we do not eliminate quadratic terms in government spending. This term is not independent of policy, as monetary policy actions affect fiscal policy decisions.

3.2 The Choice of Policy Instruments

We assume that monetary policy uses short-term nominal interest rate as the policy instrument and that the fiscal authority uses government spending as its instrument. This follows current convention for the monetary policymaker, but the choice of fiscal instrument is more arbitrary. Our choice is determined by the following considerations. First, there is no well-established form of fiscal policy rule. Empirical estimates of fiscal policy reaction functions (see Taylor 2000, Auerbach 2003, and Favero and Monacelli 2005, for example) suggest that both government spending and taxes are varied by the fiscal authority. Second, from a methodological point of view, using spending is

[9]This derivation follows Woodford (2003a). Alternative ways of deriving social welfare (see, for example, Sutherland 2002 and Benigno and Woodford 2004) are inappropriate for discretionary policy.

a convenient starting point. We discuss in section 7 how our results change if we use taxes instead, so, effectively, we look at both possibilities. We shall refer to the income tax rate as "taxes" in what follows.

3.3 The Benchmark Ramsey Allocation

The Ramsey allocation takes into account the presence of distortions, as summarized by the implementability constraints (10)–(11), the resource constraint (12), and the financial constraint (13). Specifically, the Ramsey allocation in the LQ framework solves

$$\min_{\{i_t, g_t, \tau_t\}} \frac{1}{2} \mathcal{E}_t \sum_{s=t}^{\infty} \beta^{s-t} W_s^S$$

subject to constraints (10)–(13) for all $t \geq 0$.

The Ramsey allocation requires commitment to policies and full cooperation between the authorities. In what follows we term this the *commitment solution*. We use the commitment solution as the benchmark case.

3.4 Discretionary Policy and Non-Cooperative Regimes

We assume that both monetary and fiscal authorities act non-cooperatively in order to stabilize the economy against shocks. Both authorities are assumed to act under discretion. We assume that the monetary authority chooses the interest rate to minimize the welfare loss with per-period metrics in the form of (15) subject to the system (10)–(13), and the benevolent fiscal authority chooses spending to minimize the welfare loss (15) subject to the same system.

Of course, if both authorities are benevolent, then they both use the per-period social loss as their objective function. If the monetary authority is inflation conservative, then it is assumed to deviate from the microfounded weight on inflation variability, so its per-period objective becomes

$$W_s^M = \rho_c c_s^2 + \rho_g g_s^2 + \rho_y y_s^2 + \rho_\pi \pi_s^2 + \mathcal{O}(\|\xi\|^3), \tag{16}$$

where $\rho_\pi \geq 1$.

Each of the policymakers solves an optimization problem every period. The resulting optimal policy reactions lead to stochastic equilibria that should be compared across a suitable metric. The obvious choice is the microfounded social loss.

The sequence of events and actions within a period is as follows. At the beginning of every period t, the state $\tilde{b}_t$ is known and shock η_t realizes. Then the two policymakers choose the value of i_t, g_t, and τ_t. There is a particular order of moves, and we shall study all possibilities: (i) one of the policymakers moves first, and the leader chooses the best point on the follower's reaction function or (ii) both policymakers move simultaneously and neither is able to exploit the reaction function of the other policymaker. In all cases, the policymakers know the state $\tilde{b}_t$ and take the process by which private agents behave as given. After the policymakers have moved, in the next stage the private sector simultaneously adjusts its choice variables π_t and c_t. The optimal π_t, c_t and policy i_t, g_t, and τ_t result in the new level of $\tilde{b}_{t+1}$ by the beginning of the next period $t+1$.

4. Calibration

We take the model's frequency to be quarterly. To achieve a steady-state rate of interest of approximately 4 percent, we set the household discount rate β to 0.99. The remaining parameters of the utility function are typical of those used in the literature; see, e.g., Canzoneri et al. (2006). The elasticity of intertemporal substitution σ is taken as 1/1.5, the Calvo parameter γ is set at 0.75 so as to imply average contracts of about a year, the elasticity of demand is taken as $\varepsilon = 7.0$ to achieve a 17 percent markup, and elasticity of labor demand is taken as $\psi = 1/3$.

The ratio of government consumption to output in the point of linearization, $1-\theta$, is a function of relative weight on the utility-from-public-consumption term, ξ. We choose to calibrate ξ such that it would result in realistic $1-\theta = 0.25$.

We shall demonstrate that different values of χ can result in qualitatively different policy interactions. In what follows, we consider two values for the debt-to-GDP ratio. Our "high" debt level, $\chi = 1.2$, corresponds to 30 percent of annual output, which is still less than the level of debt in a number of European economies. However, we only consider one-period debt, so the figure of 30 percent is large enough to demonstrate *qualitative* difference with $\chi = 0$, which corresponds to 0 percent of annual output and which we treat as "low" debt level.

The steady-state level of tax rate is a function of the steady-state debt level, $\tau = \tau(\chi)$. Of course, "high" and "low" level of debt corresponds to "low" and "high" tax rate.

We only study the effect of cost-push shocks, as they are the biggest potential source of social loss.[10] We calibrate the standard deviation of an i.i.d. cost-push shock as 0.005. In our baseline case, this generates a standard deviation for inflation of 0.0038, approximately the same order of magnitude as experienced in developed countries over the last couple of decades. This number is also not unreasonable given other academic studies. Ireland (2004) uses a cost-push shock, which is AR(1) with a standard deviation of 0.0044. Smets and Wouters (2003) report an i.i.d. cost-push shock with a much smaller standard deviation in the model with inflation persistence, while Rudebusch (2002) estimates a standard deviation of 0.01 for an i.i.d. cost-push shock.

5. Benevolent Policymakers

We first look at how monetary and fiscal policy interact if policy objectives are not distorted. As we shall see, the qualitative results crucially depend on calibration of one particular parameter, the steady-state level of debt. We find it convenient to emphasize this difference from the very beginning, but we defer the remaining robustness analysis to section 7.

We assume that both authorities have *identical* intraperiod *social* objectives (15). We solve the optimization problem and find optimal policy and the incurred costs for jointly optimal policies under discretion and under commitment. Figure 1 plots the impulse responses to a unit cost-push shock for the low- and high-debt cases.[11] We report and discuss several results.

5.1 Commitment vs. Discretion

Optimal policy under commitment, compared with discretionary policy, requires smaller and smoother movements in policy instruments.

[10] We have examined what happens if we introduce taste or productivity shocks, and all quantitative results are virtually the same.

[11] We only show several initial periods. With time, all variables (including debt) converge back to the steady state.

Figure 1. Impulse Responses to a Unit Cost-Push Shock under Cooperative Optimization of Benevolent Authorities

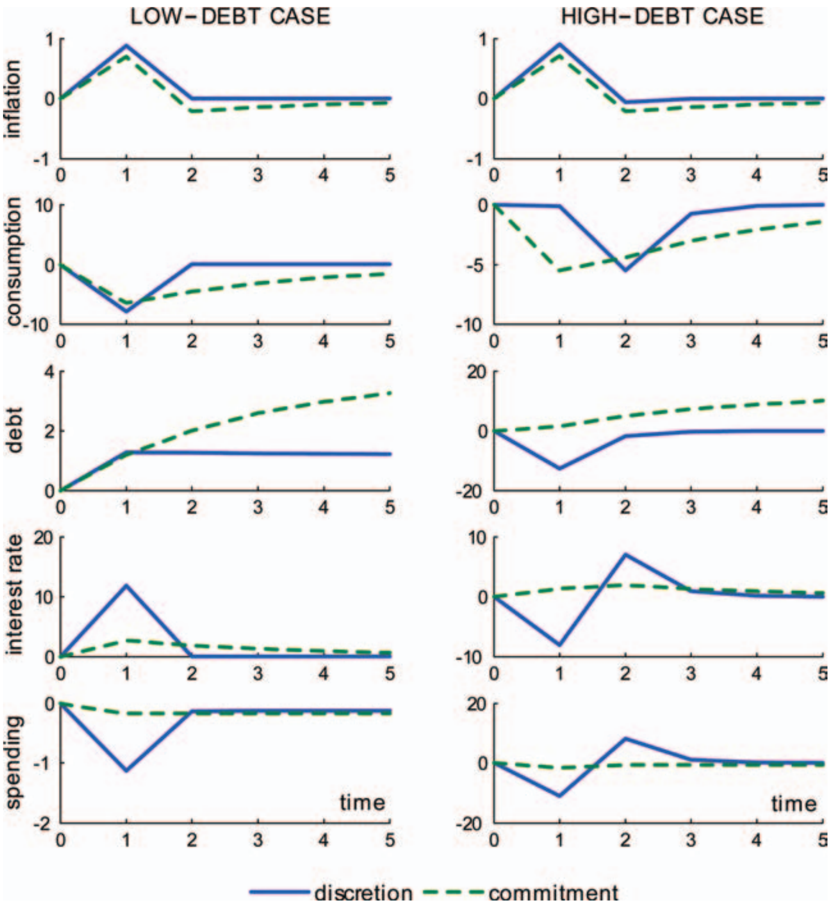

Note: Fiscal policy uses spending as an instrument.

Commitment policy has control over all future states, including expectations. The policy is chosen once and the policymaker is able to adhere to it. As a result, the private sector sets expectations that help stabilization. By promising to keep the interest rate above the baseline in all future periods, the monetary policymaker is able to achieve an immediate fall in inflation and overall price stability. Also, only a small initial rise in the interest rate is sufficient for this stabilization.

In contrast, the discretionary policymaker reoptimizes every period. The discretionary policymaker cannot promise to keep the interest rate above the baseline in future periods, as the private sector knows that all future policymakers will reoptimize. The policymaker, therefore, has no control over private-sector expectations along the whole future paths. It has to raise the interest rate by much more in the first period in order to move all variables sufficiently close to the steady state already after the first period.[12] This distinction in qualitative behavior of commitment and discretionary policies is well known and is frequently exploited in the design of delegation schemes—different from the one we consider here—that make discretionary policy smooth to replicate commitment policy and achieve higher social welfare.[13]

5.2 Non-Cooperative Discretionary Policy Regimes

All non-cooperative regimes with identical objectives deliver the same equilibrium. This result shows that strategic behavior can only be of importance if there is conflict in objectives. That the order of moves is of no importance is a direct consequence of policymakers having identical objectives *given* that the equilibrium is unique. In this setting, neither authority is trying to exploit the other in a pursuit of their own target—they internalize externalities.[14]

This result holds because there is a unique discretionary equilibrium. By uniqueness we mean that following a shock, the economy

[12]The inability of discretionary policy to control expectations of the private sector about future policy and future states also implies that the choice of policy instruments might matter. In contrast to Benigno and Woodford (2004), we cannot use consumption or inflation as a policy instrument, as the discretionary policymaker has no control of their future values.

[13]A non-exhaustive list includes Svensson (1997), Walsh (2003), Woodford (2003b), and Vestin (2006).

[14]Formally, it is straightforward to demonstrate that a Nash equilibrium with identical objectives coincides with the cooperative equilibrium. The system of first-order conditions for a single optimization problem with two instruments in cooperative equilibrium will be identical to the two systems of first-order conditions for the two separate maximization problems for each of the instruments. Similarly we can deal with Stackelberg equilibrium: the systems of first-order conditions will contain some additional terms, but they are all zero if the objectives are the same.

follows a unique path which satisfies the conditions of time consistency and optimality; see appendix 1. This is not an obvious result a priori given that a similar model with fiscal policy, which controls debt by means of an *exogenous* feedback-on-debt rule, is shown to have multiple discretionary equilibria in Blake and Kirsanova (2008). In that model, if the fiscal feedback on debt is sufficiently strong, so any debt displacement is corrected by fiscal means very quickly, monetary policy behaves in a conventional way. In response to a cost-push shock, monetary policy raises the interest rate, which leads to an initial increase in debt and also slows down the debt stabilization process. The strong fiscal feedback ensures the debt stability. If fiscal feedback is zero, so fiscal policy does not control debt at all, monetary policy lowers the interest rate in response to a cost-push shock. As a result, the lower real rate reduces the level of debt below the steady state. The subsequent increase in the interest rate reduces output and inflation and also pushes debt back to the steady state. There is an intermediate case where fiscal policy controls debt only weakly, so both monetary regimes are possible.

The multiplicity of regimes is possible because of the complementarity between the decision variable of the private sector (inflation) and the instrument of the monetary policymaker (the interest rate) in their effects on debt, as defined in Cooper and John (1988). Namely, a too-low debt stock can be increased with either a high interest rate or with low inflation, and, crucially, an increase in the interest rate reduces inflation. This complementarity exists in our model too, but its strength depends on fiscal stance: when fiscal feedback is sufficiently strong, then debt is stabilized by fiscal means and effects of inflation and interest rate on the speed of debt stabilization are negligible. In contrast to Blake and Kirsanova (2008), the fiscal authority is a strategic player in our model in this paper, and it optimally chooses to reduce spending in a response to a higher interest rate, i.e., chooses sufficiently strong fiscal feedback to put the economy outside the area of multiple equilibria.[15]

[15] It is easy to prove analytically that the discretionary equilibrium is unique in the "low-debt" case, as algebra can be substantially simplified as in Blake and Kirsanova (2008).

5.3 The Level and Dynamics of Government Debt

The policy mix very much depends on the *steady-state level* of debt, χ. The substantial difference is in terms of the directions of optimal responses to the same shock.

In the low-debt case, contractionary monetary policy is supported by contractionary fiscal policy in a response to a positive cost-push shock. This leads to a rise in debt. Fiscal policy is then used to bring debt back to the steady state: government spending stays at a reduced level for a long time.

In the high-debt case, the interest rate *falls* in the first period in a response to a positive cost-push shock. As the first-order effect of interest rates on debt is large, a fall in interest rates reduces the level of domestic debt. Moreover, the fiscal contraction reduces debt even further. This first-period response allows monetary policy to raise the interest rate in the second period and reduce inflation. (Note that *second-period* inflation overshoots the steady-state level. This helps to reduce inflation in the *first* period too, as the rational private sector sets prices lower in anticipation of this lower future price.) Fiscal policy also raises spending, as this not only stabilizes debt but also reduces recession caused by high interest rates.

The left panel in figure 2 demonstrates that the social loss is a non-monotonic function of χ. This is consistent with a striking change in the way the stabilization policy works. When χ is below some threshold level, then any further rise in χ creates more problems for policymakers if monetary policy raises the interest rate in response to a cost-push shock. A high interest rate contributes to debt accumulation, and the effect is proportional to χ. When χ is below this threshold, the gains from inflation stabilization in the first periods after the shock outweigh the losses from the slow stabilization of the economy because of slow debt dynamics. When χ is above this threshold, it becomes welfare improving to stabilize debt quicker. Therefore, it becomes optimal to lower the interest rate in the first period after the shock and raise it in the second period. This policy leads to faster debt stabilization and also curtails inflation. To emphasize the difference, we have chosen the low- and high-debt cases on either side of the "hump."

The *dynamic* process of debt accumulation plays a very important role for the stabilization policy mix. First, its presence imposes

Figure 2. The Effect of Steady-State Level of Debt on the Social Loss

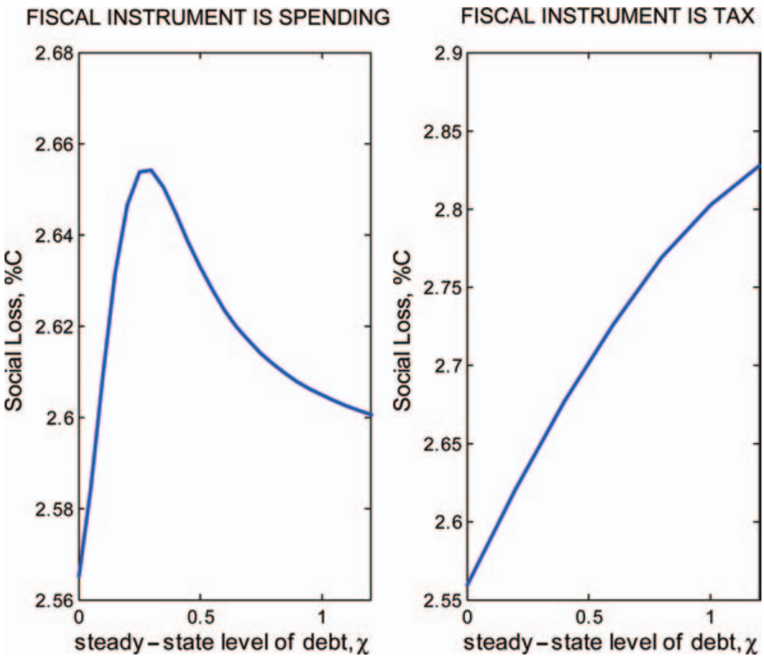

Notes: The figure shows the joint monetary-fiscal optimization of benevolent policymakers under discretion. The fiscal policymaker uses different fiscal instruments.

the requirement on the policy mix to prevent an explosion of debt. Second, its presence alters the dynamics of stabilization. If there is no debt, then following a shock a discretionary policy stabilizes the economy within *one* period. Policy can only reduce the *amplitude* of reactions to shocks within the first period. If debt accumulation is present, then policy can also reduce the half-life of effects of shocks that have entered the system. In other words, it enables stabilization to be smoothed over many periods, which may or may not be welfare improving. With the presence of debt, the private sector's expectations affect the economy for more than one period, as the evolution of debt can be affected by the forecast of future policy. Expectations set in period t affect the whole future path of state variables and they affect policy decisions taken in period $t + k$, $k \geq 1$. The way

Table 1. Social Loss for Different Scenarios with Different Fiscal Instruments, % of Steady-State Consumption

	Low-Debt Case			High-Debt Case		
Fiscal Instrument →	**G**	**T**	**G&T**	**G**	**T**	**G&T**
Commitment	2.03	0.23	0.22	2.05	0.23	0.22
Cooperation under Discretion	2.57	2.56	2.54	2.60	2.83	2.73
Stabilization Bias	0.54	2.34	2.33	0.55	2.60	2.51

monetary policy stabilizes inflation in the case of high debt is an example of how expectations of future policy can be exploited.

The stabilization bias is around 0.5 percent of the steady-state consumption level that should be given up in order to compensate for reduced volatility. We compute the stabilization bias as a difference between the social loss under discretion and the social loss under commitment (Ramsey allocation). Table 1 presents more detailed information. Together with the left panel in figure 2, the table suggests that the level of debt can affect the size of stabilization bias, and the order of the effect can be around 0.1 percent of the steady-state consumption level, which is a relatively big number. We shall use these numbers as a benchmark. All our future computations of losses will be relative to the loss under the *benevolent discretionary* policy, and to obtain the stabilization bias one has to add a corresponding number from table 1.[16]

This section has established that with benevolent policymakers, the order of moves and information structure is of no consequence. However, there are issues that arise from the dynamic nature of the model through the steady-state level of debt. In what follows we need to investigate the importance of steady-state levels of debt together with the conservative central bank proposal.

6. Conservative Central Bank

Suppose the fiscal authorities are benevolent, but we allow the monetary authorities to place a higher relative weight on inflation

[16]Here and everywhere else we measure the loss in percentage of steady-state consumption level.

Figure 3. Social Welfare Loss as a Function of Monetary Conservatism for Different Non-Cooperative Regimes

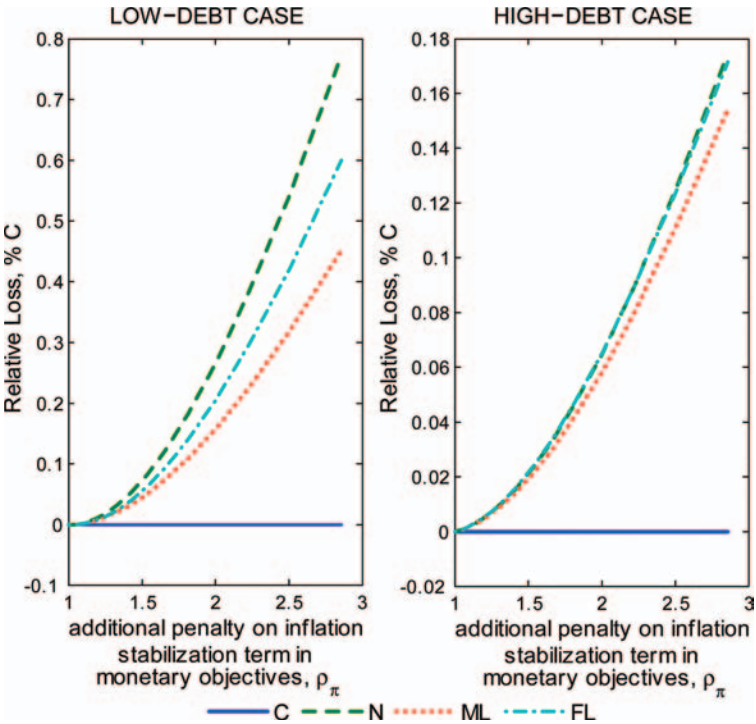

Note: Fiscal policy uses spending as an instrument.

stabilization, by adopting objective (16) with $\rho_\pi > 1$. How does this affect social welfare? We keep the fiscal authorities benevolent.

The left panel of figure 3 suggests that the loss quickly rises with the degree of monetary conservatism for all three non-cooperative regimes; i.e., we can double the stabilization bias very quickly. When the penalty is very close to one, and in the low-debt case, there is an extremely small social gain for both regimes of fiscal leadership and of simultaneous moves, but the monetary leadership regime is unambiguously worse than the cooperation of benevolent authorities. If the steady-state level of debt is large, then there is no social gain for any degree of inflation conservatism and the social loss rises quickly with the degree of inflation conservatism. To understand these results, it is instructive to look at impulse responses to a unit

Figure 4. Impulse Responses to a Unit Cost-Push Shock for Different Strategic Regimes if the Monetary Authority Is Conservative with Small Degree of Conservatism, $\rho_\pi = 0.06$

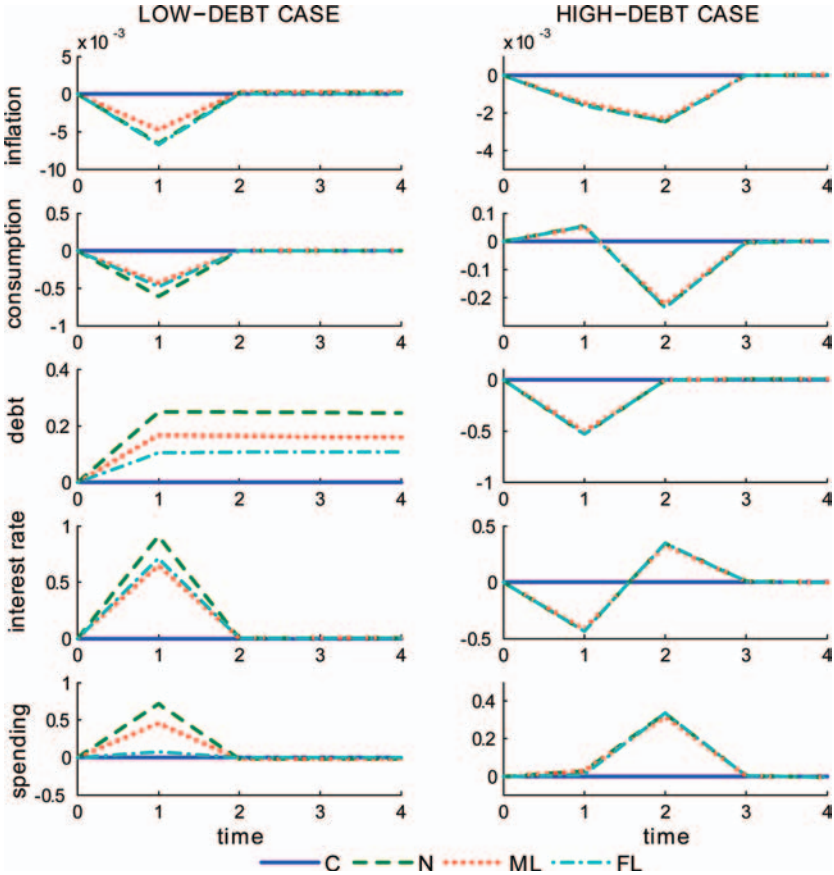

Notes: Fiscal policy uses spending as an instrument. All responses are shown as relative to those under cooperation.

cost-push shock in figure 4. We plot differences with the cooperative solution, which itself is plotted in figure 1.

6.1 *Simultaneous Policy Decisions*

In the *low-debt* case, the monetary authority reacts more actively to a cost-push shock than if it were benevolent. It is more concerned with inflation variability than society and is prepared to pay for gains with higher variability of the demand-related components. The monetary authority, therefore, raises the interest rate higher to eliminate inflation more aggressively. The benevolent fiscal authority tries to eliminate the resulting recession. It therefore contracts less (expands more) than if both authorities were benevolent, although inflation is still reduced. The reduction in the cost of inflation volatility only outweighs the cost of fiscal volatility if the degree of monetary conservatism is extremely small. With a higher degree of conservatism of the monetary authorities, the implied fiscal volatility becomes very costly.

In the *high-debt* case, it is optimal for the monetary authority to take into account the debt stabilization issues. It chooses to reduce the interest rate in the first period by more, but the resulting smaller fiscal contraction allows the monetary authorities to then raise the interest rate by more in the second period without having an adverse effect on debt. Inflation, thus, overshoots more (falls relative to the cooperative benevolent case) in the second period. It, therefore, rises less (or falls relative to the cooperative case) in the first period. Again, the reason for its reduction in the first period is the second-period overshooting and the rational expectations of price setters: A rational private sector perceives the second-period fall of prices and sets prices low in the first period.

The simultaneous-move regime leads to more aggressive policy than under cooperation, and this results in lower inflation but also in higher volatility for demand-related terms and instruments, and therefore, in a more costly equilibrium.

6.2 Fiscal Leadership

In the *low-debt* case, the fiscal authority knows that if it contracts less than in the cooperative scenario, then this will cause the monetary authority to contract more in order to fight excess inflation. Moreover, the monetary authority will contract even more due to its inherent conservatism. So the fiscal authority chooses to contract only slightly less than in the cooperative scenario. The monetary authority still contracts *more* than in the cooperative scenario, and overall this results in slightly lower inflation. Inflation falls nearly as much as under Nash, because fiscal authorities do not expand as much as under Nash. Debt rises only slightly higher than in the cooperative scenario, and less than if both authorities move simultaneously. With a higher degree of monetary conservatism, the monetary authorities have to pay too high an output and consumption cost for lower inflation, so the regime is unambiguously worse than full cooperation.

In the *high-debt* case, the authorities have more problems with debt stabilization. The fiscal authority contracts more in the first period, thus not allowing monetary policy to expand as much as under Nash. The monetary authority therefore contracts less in the second period. This creates less inflation overshooting than under Nash. The welfare loss is bigger than under Nash for a moderate degree of monetary conservatism. However, with a sufficiently high degree of conservatism, the monetary-fiscal interaction results in smaller consumption volatility than under Nash, and this improves welfare slightly. The regime becomes slightly better than the simultaneous-move regime, although the quantitative response is extremely close to it. The debt stabilization task dominates the concerns of both authorities and nearly equally constrains instrument movements in all three non-cooperative regimes.

6.3 Monetary Leadership

In the *low-debt* case, the leading monetary authority knows that raising the interest rate causes the fiscal authority to try to offset the resulting recession. This consideration would stop the *benevolent* monetary authority from raising interest rates. However, as the monetary authority is conservative, the cost of moving g becomes *relatively* less important, so it does raise the interest rate. But it raises the interest rate less than in the simultaneous-move regime. Lower pressure on debt allows fiscal policy to offset the effect more efficiently. Inflation falls less than under Nash and this determines the loss.

In the *high-debt* case, the monetary authority contracts more in the first period than under Nash and the fiscal authorities have to reduce spending by more. Debt itself does not fall as much as it does under Nash. The monetary authority is unable to achieve as large an inflation overshoot in the second period, and the consequent overall gain in inflation stabilization, as under Nash. It is able, however,

Figure 5. Social Loss as a Function of Conservatism for Different Strategic Regimes and Policy Arrangements

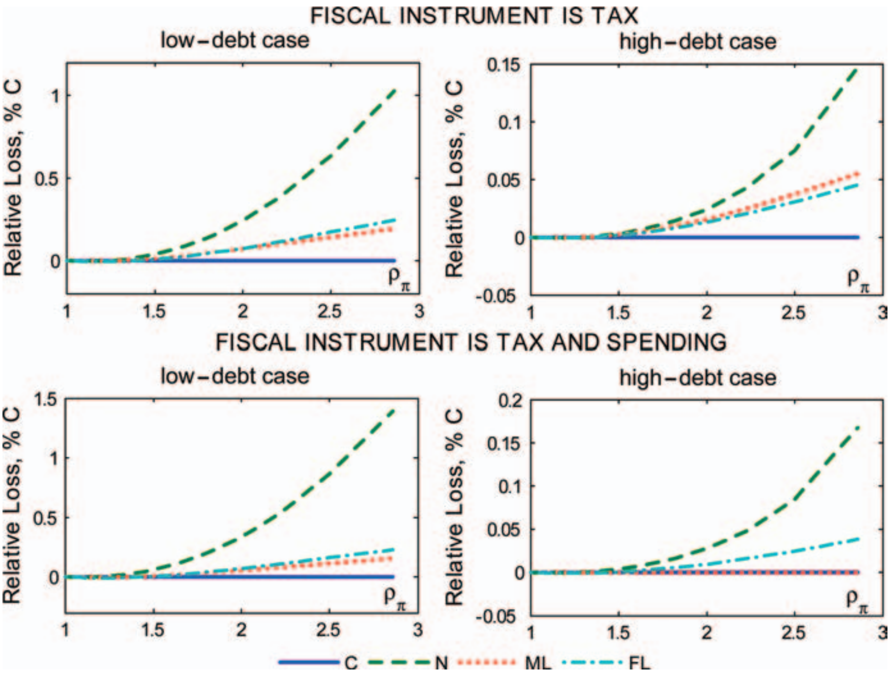

Note: Fiscal policy uses taxes as an instrument or it uses both instruments, taxes and spending.

to achieve smaller variability of demand-related terms than under Nash. Both effects taken together result in a smaller welfare loss than for the simultaneous-move case.

7. Robustness: Using Tax as an Instrument

7.1 Tax as Fiscal Instrument

We have shown how our results change with the level of steady-state debt. It is also important to look at the choice of fiscal instrument. We rerun our simulations assuming that (i) fiscal policy uses tax as the instrument and (ii) fiscal policy uses two instruments, both income tax rate and spending. Figure 5, which repeats figure 3 but is plotted for taxes (the upper panels) and taxes and spending (the lower panels), suggests that our main conclusion remains valid: generally speaking, delegating monetary policy to a conservative central bank does not improve social welfare.

There are some important differences with the case where only spending is used. Consider using tax as a single fiscal instrument and suppose that authorities are *benevolent*. As taxes affect marginal cost directly, fiscal policy tries to offset any cost-push shocks immediately by lowering the tax rate. Indeed, the first column of plots in figure 6 suggests that in response to a unit cost-push shock, taxes fall, and this allows monetary policy to raise the interest rate without any solvency concerns.[17] The debt is also stabilized *by taxes*: the first-period reduction in the tax rate and high interest rates require higher taxes in subsequent periods; taxes rise and bring debt back to the steady state very quickly. As a result, the dynamics of debt is very different from the one in the case where fiscal policy can only use spending as an instrument: debt quickly converges to the steady state. Taxes have no direct effect on demand and consumption, so fiscal policy can eliminate debt displacement in consequent periods more efficiently with fewer externalities, and monetary policy can stabilize inflation.

Interestingly, the welfare gain from using taxes instead of spending is very small: the loss under the benevolent discretionary monetary-fiscal policy in the low-debt case is only 0.006 percent higher if fiscal policy uses spending instead of taxes as the only instrument. Moreover, in the high-debt case, using taxes is more costly than using spending. To understand these results, it is instructive to compare them with those under commitment of benevolent authorities which would deliver the highest welfare in the absence of dynamic lump-sum taxes. Benigno and Woodford (2004) demonstrate that fiscal policy under commitment is very successful in eliminating all effects of a cost-push shock on the economy. These results are consistent with ours; see the first line in table 1: the loss under the commitment policy falls substantially if taxes are used. Under

[17]We have checked that impulse responses are qualitatively similar for all degrees of conservatism and for low- and high-debt scenarios. So we only present small conservatism and a low-debt case in figure 6.

Figure 6. Impulse Responses to a Unit Cost-Push Shock When Fiscal Policy Uses Different Instruments

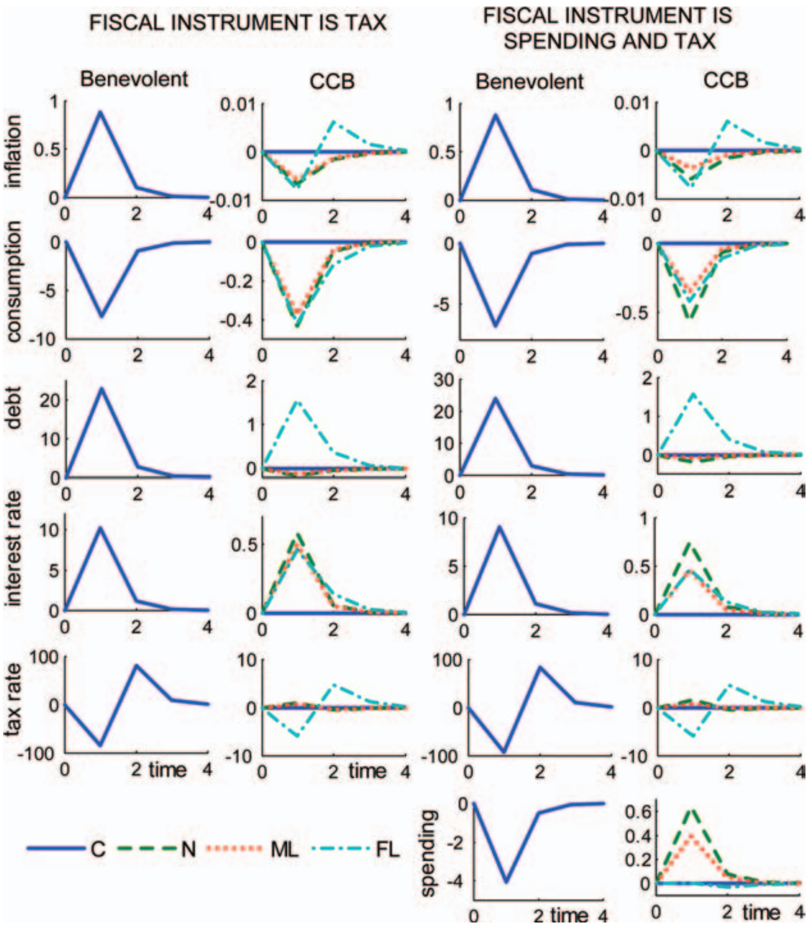

Notes: Columns 2 and 4 plot relative responses, relative to the case of cooperation of benevolent authorities. The "low-debt" scenario and small degree of conservatism, $\rho_\pi = 0.06$, are assumed in all cases where applicable.

discretion, however, taxes cannot move as freely as under commitment, and they cannot efficiently offset the effect of cost-push shocks, although they still move in the right direction. The stabilization bias is large, around 2.5 percent, only because the commitment policy is so successful.

The right panel in figure 2 suggests that with a higher level of steady-state debt, χ, the social loss rises. The level of χ determines the size of the first-order effect of the interest rate on debt accumulation. With higher χ, the problem of debt stabilization becomes more difficult. Taxes become predominantly used in the debt stabilization task; they do not fall by much, and this makes the task of inflation stabilization more difficult for monetary policy.

The relative ranking of different leadership regimes barely changes. The simultaneous-move regime leads to most welfare losses. However, it is not the most aggressive regime. The second column of plots in figure 6 illustrates this. As before, we plot impulse responses relative to those under cooperation of benevolent policymakers. As all qualitative responses are very similar for different degrees of conservatism and for the low- and high-debt scenarios, we only plot the case with low conservatism and with a low level of debt.

The second column of plots in figure 6 demonstrates interactions for different leadership regimes in the case of a *conservative central bank*. Under the *Nash regime*, the monetary authority raises the interest rate more than if it was benevolent. It also reduces inflation by more and reduces consumption. The optimal response of the fiscal authority is to raise taxes. If monetary policy is not taken into account, then higher taxes would allow the cost-push shock to have a bigger effect on inflation, the real interest rate would fall, and consumption would rise. In the simultaneous-move regime, the fiscal authority, therefore, reduces taxes by less than it does under cooperation. Monetary policy raises the interest rate higher and so on. In equilibrium, inflation is reduced but consumption falls. For a small degree of conservatism, the gain of lower inflation outweighs the loss of lower consumption and output, but consumption volatility rises very fast with the degree of monetary conservatism, and the regime quickly becomes welfare inferior.

If the *monetary authority leads*, it knows that the fiscal authority will not reduce taxes as if it were benevolent, so it raises the interest rate by less and the fiscal authority reduces tax by less than under Nash. This leads to nearly the same outcome as under Nash if the degree of conservatism is small. With a higher degree of conservatism, policy aggressiveness leads to more volatility in consumption, which is still less than under Nash. So the monetary leadership regime is preferable over the Nash regime.

If the *fiscal authority leads*, it knows that the monetary authority will raise interest rates, so it consequently reduces taxes by more. It explicitly exploits the monetary policy reaction function: the monetary authority raises the interest rate, but by less then if it were benevolent because the fiscal authority does part of the disinflation. As a result of these first-period movements, debt rises higher. Fiscal policy thus has to respond by raising taxes in the second period, which also results in higher inflation in the second period. If the degree of monetary conservatism is small, then the fiscal leadership regime delivers the lowest loss among the three non-cooperative regimes, because of the relatively large fall in inflation. The second-period inflation hike, however, dominates the large social welfare loss if the monetary authority has a higher degree of conservatism.

Neither of the non-cooperative regimes can be defined as the most aggressive here: the Nash regime results in higher variability in the interest rate while the fiscal leadership regime results in the biggest volatility of the fiscal instrument.

7.2 Tax and Spending as Fiscal Instruments

Adding government spending as the second fiscal instrument does not change any of the welfare results, as the two lower plots in figure 5 demonstrate. An inflation-conservative monetary authority, generally speaking, generates social welfare loss. If the authorities are benevolent, then the social loss is only marginally smaller than the loss from stabilization if only taxes are used; see table 1.

The similarity is also apparent from impulse responses shown in the last two columns of plots in figure 6. As debt is now stabilized by taxes, spending is optimally chosen to help monetary policy to reduce inflation if the authorities are benevolent. If there are distorted objectives, then interest rate/tax interactions are nearly the same as if tax were the only fiscal instrument. The possibility to use spending as well makes little difference: spending plays the same role (relative to the case of cooperation of benevolent policymakers) as in the scenario where it was the only fiscal instrument; one can compare the last column of plots in figure 6 with the first column of plots in figure 4.

8. Summary of Results and Conclusions

This paper presents a detailed account of discretionary monetary and fiscal policy interactions in a fully specified intertemporal general equilibrium model with particular emphasis on non-cooperative interactions under inflation conservatism of the monetary authority.

We demonstrate that if the fiscal policymaker is benevolent but acts strategically, then delegating monetary policy to a policymaker that puts a higher than socially optimal weight on inflation stabilization generally increases the stabilization bias. Such a distortion of the otherwise social objectives of one policymaker brings the two policymakers into conflict with each other, and this is welfare destructive. This example demonstrates that what works well in an economy with one strategic policymaker may not work at all in an economy with two strategic policymakers, as the second strategic policymaker can offset all actions of the first one and vice versa.

Of course, this result does not imply that the problem of optimal delegation has no solution if there are several strategic policymakers. We have studied only one particular delegation scheme: the conservative central bank scenario. We have distorted the relative weights on social objectives, but we did not introduce any additional objectives in order to have a simple and clear experiment. Additional non-microfounded terms in policymakers' objectives (like instrument costs) or the use of different policy instruments (like VAT or a consumption tax) might have different effects on the ability and willingness of the policymakers to engage in conflict with each other. Different degrees of precommitment may affect the degree of conflict too. We leave these and similar issues for future research. Our result, however, suggests that making fiscal authorities flexible and strategic may have costs, and these should be taken into account. We also show how conflict arises if the weights on social objectives are changed but all objectives remain social.

To arrive at our main conclusions, we have investigated monetary and fiscal policy interactions under discretion in a fairly standard model with an explicit budget constraint for the fiscal authority. The analysis also yields some additional insights.

First, the choice of fiscal instrument is important, although not for assessing gains from delegation where there are losses whatever the fiscal instrument. The transmission mechanism differs considerably: (i) Taxes are most useful in stabilizing debt, and the way they are optimally used does not depend on the level of steady-state debt, and (ii) spending has more limited powers to stabilize debt, but has a large effect on domestic demand. If spending is the only fiscal instrument, then the size of steady-state debt has an important qualitative effect on policy interactions. If the steady-state level of debt is high, then monetary policy has to take an active part in debt stabilization.

Second, among the three non-cooperative regimes, the Nash one is unlikely to be welfare dominating. In most cases, this regime leads to large movements of the policy instruments, which typically implies increased volatility of the key economic variables. Most of our results suggest that monetary leadership is relatively better if the monetary authority is inflation conservative, but this hinges on the social welfare metric and requires that inflation stabilization has far greater weight than any other target variable.

Finally, this paper offers an additional contribution to the literature: we offer a methodological approach to solving a non-cooperative leadership equilibria in an LQ RE model with two policymakers. This approach can be easily modified to study different types of equilibria and interactions of many agents, say in a multi-country setting.

Appendix 1. Leadership Equilibria under Discretion

This section demonstrates how to solve non-cooperative dynamic games in the linear-quadratic rational expectations framework. Our definition of discretionary policy is conventional and is widely used in the monetary policy literature; see, e.g., Oudiz and Sachs (1985), Backus and Driffill (1986), Clarida, Galí, and Gertler (1999), and Woodford (2003a). Currie and Levine (1993) demonstrate how to solve Nash games. We therefore only describe leadership equilibria.[18]

Discretionary Policy

We assume a non-singular linear deterministic rational expectations model, augmented by a vector of control instruments.[19]

[18] A simultaneous-move regime will be a particular case of the leadership regime.

[19] None of the results presented here depend on the deterministic setup outlined and the consequent assumption of perfect foresight. Shocks can be included into vector y_t; see, e.g., Anderson et al. (1996) and Blake and Kirsanova (2008).

Specifically, the evolution of the economy is explained by the linear system

$$\begin{bmatrix} \mathsf{y}_{t+1} \\ \mathsf{x}_{t+1} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} \mathsf{y}_t \\ \mathsf{x}_t \end{bmatrix} + \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} \begin{bmatrix} \mathsf{u}_t^L \\ \mathsf{u}_t^F \end{bmatrix}, \tag{17}$$

where y_t is an n_1-vector of predetermined variables with initial conditions y_0 given, x_t is an n_2-vector of non-predetermined (or jump) variables with $\lim_{t\to\infty} \mathsf{x}_t = 0$, and u_t^F and u_t^L are the two vectors of policy instruments of two policymakers, named F and L, of size k_F and k_L, respectively. For notational convenience we define the n-vector $\mathsf{z}_t = (\mathsf{y}_t', \mathsf{x}_t')'$, where $n = n_1 + n_2$, and the k-vector of control variables $\mathsf{u}_t = (\mathsf{u}_t^{L\prime}, \mathsf{u}_t^{F\prime})'$, where $k = k_F + k_L$. We assume the equations are ordered so that A_{22} is non-singular.

Typically, the second block of equations in this system represents an aggregation of the first-order conditions to the optimization problem of the private sector, which has decision variables x_t. Additionally, there is a first block of equations which explains the evolution of the predetermined state variables y_t. These two blocks describe the "evolution of the economy" as observed by policymakers.

The intertemporal welfare criterion of policymaker i, $i \in \{L, F\}$, is defined by the quadratic loss function

$$W_t^i = \frac{1}{2}\sum_{s=t}^{\infty} \beta^{s-t}\left(g_s^{i\prime} \mathcal{Q}^i g_s^i\right) = \frac{1}{2}\sum_{s=t}^{\infty} \beta^{s-t}\left(\mathsf{z}_s' Q^i \mathsf{z}_s + 2\mathsf{z}_s' P^i \mathsf{u}_s + \mathsf{u}_s' R^i \mathsf{u}_s\right). \tag{18}$$

The elements of vector g_s^i are the goal variables of policymaker i, $g_s^i = \mathcal{C}^i(\mathsf{z}_s', \mathsf{u}_s')'$. Matrix $\mathcal{Q}^i$ is assumed to be symmetric and positive semi-definite.

Assumption 1. *Suppose that at any time t, the private sector and policymakers respond only to the current state*

$$\mathsf{u}_t^L = \mathcal{F}^L(\mathsf{y}_t) = -F^L \mathsf{y}_t, \tag{19}$$

$$\mathsf{u}_t^F = \mathcal{F}^F(\mathsf{y}_t) = -F^F \mathsf{y}_t, \tag{20}$$

$$\mathsf{x}_t = \mathcal{N}(\mathsf{y}_t) = -N \mathsf{y}_t. \tag{21}$$

This assumption rules out non-stationarity of policy and private-sector decisions—i.e., any time dependence from the more general formulation $\mathsf{u}_t^i = \mathcal{F}^i(t; \mathsf{y}_t, \mathsf{y}_{t-1}, .., \mathsf{y}_{t-k}, ..)$, $\mathsf{x}_t = \mathcal{N}(t; \mathsf{y}_t, \mathsf{y}_{t-1}, .., \mathsf{y}_{t-k}, ..)$—and restricts policy decisions to memoryless feedback rules. We also assume that rules are *linear* in the state.

We define discretionary policy as satisfying several constraints. We want to assume that the policymaker can implement at each point of time its policy decision *before* the private sector selects its own action x_t.

ASSUMPTION 2. *At each time t, the private sector observes the current decision u_t and expects that future policymakers at any time $s > t$ will reoptimize, will apply the same decision process, and implement policy (19)–(20).*

PROPOSITION 1. *Given assumption 2, the current aggregate decision of the private sector can be written as a linear feedback function*

$$\mathsf{x}_t = -J\mathsf{y}_t - K\mathsf{u}_t = -J\mathsf{y}_t - K^F\mathsf{u}_t^F - K^L\mathsf{u}_t^L, \tag{22}$$

where

$$J = (A_{22} + NA_{12})^{-1}(A_{21} + NA_{11}), \tag{23}$$

$$K^F = (A_{22} + NA_{12})^{-1}(B_{22} + NB_{12}), \tag{24}$$

$$K^L = (A_{22} + NA_{12})^{-1}(B_{21} + NB_{11}). \tag{25}$$

Proof. Relationship (21) can be taken with one lead forward, and y_{t+1} is substituted from the first equation (17). We obtain

$$\begin{aligned} \mathsf{x}_{t+1} &= -N\mathsf{y}_{t+1} = -N(A_{11}\mathsf{y}_t + A_{12}\mathsf{x}_t + B_{11}\mathsf{u}_t^L + B_{12}\mathsf{u}_t^F) \\ &= A_{21}\mathsf{y}_t + A_{22}\mathsf{x}_t + B_{21}\mathsf{u}_t^L + B_{22}\mathsf{u}_t^F, \end{aligned}$$

from where it follows

$$\begin{aligned} \mathsf{x}_t = &-(A_{22} + NA_{12})^{-1}\big((A_{21} + NA_{11})\mathsf{y}_t + (B_{22} + NB_{12})\mathsf{u}_t^F \\ &+ (B_{21} + NB_{11})\mathsf{u}_t^L\big) = -J\mathsf{y}_t - K\mathsf{u}_t, \end{aligned}$$

where J and $K = (K^F, K^L)$ are defined as in (23)–(25). Invertibility of A_{22} ensures invertibility of $A_{22} + NA_{12}$ almost surely.

Proposition 1 implies that the policymakers, which move before the private sector, take into account their "instantaneous" influence on the choice of x_t, which is measured by $-K$.

ASSUMPTION 3. *At each time t, policymaker F knows assumptions 1 and 2 and observes the current decision u_t^L of policymaker L.*

PROPOSITION 2. *Given assumption 3, the current decision of policymaker F can be written as a linear feedback function:*

$$\mathsf{u}_t^F = -G\mathsf{y}_t - D\mathsf{u}_t^L. \tag{26}$$

Proof. We shall prove it as part of proof of proposition 3.

Proposition 2 implies that policymaker L, which moves before the private sector and before policymaker F, takes into account its "instantaneous" influence on the choice of u_t^F, which is measured by $-D$. It immediately follows that

$$F^F = G - DF^L. \tag{27}$$

ASSUMPTION 4. *At each point in time t, policymaker L knows assumptions 1, 2, and 3.*

DEFINITION 1. *Policies determined by (19)–(20) are discretionary (under intraperiod leadership of policymaker L) if each policymaker finds it optimal to continue to follow its policy in every period $s > t$, given assumptions 1–4.*

Policy of the Follower. The policy of the follower, u_t^F, satisfies the following Bellman equation:

$$V_t^F(\mathsf{y}_t) = \min_{\mathsf{u}_t^F} \left(\frac{1}{2} \left(\mathsf{y}_s' \hat{Q} \mathsf{y}_t + 2\mathsf{y}_t' \hat{P} \mathsf{u}_t + \mathsf{u}_t' \hat{R} \mathsf{u}_t \right) + \beta V_{t+1}^F (\hat{A}\mathsf{y}_t + \hat{B}\mathsf{u}_t) \right), \tag{28}$$

with

$$\hat{Q} = Q_{11}^F - Q_{12}^F J - J' Q_{21}^F + J' Q_{22}^F J,$$
$$\hat{P}_1 = J' Q_{22}^F K^L - Q_{12}^F K^L + P_{11}^F - J' P_{21}^F,$$

$$\begin{aligned}
\hat{P}_2 &= J'Q^F_{22}K^F - Q^F_{12}K^F + P^F_{12} - J'P^F_{22},\\
\hat{R}_{11} &= K^{L\prime}Q^F_{22}K^L - K^{L\prime}P^F_{21} - P^{F\prime}_{21}K^L + R^F_{11},\\
\hat{R}_{12} &= K^{L\prime}Q^F_{22}K^F - K^{L\prime}P^F_{22} - P^{F\prime}_{21}K^F + R^F_{12},\\
\hat{R}_{22} &= K^{F\prime}Q^F_{22}K^F + R^F_{22} - K^{F\prime}P_{22} - P^{F\prime}_{22}K^F,\\
\hat{A} &= A_{11} - A_{12}J,\\
\hat{A}_1 &= B_{11} - A_{12}K^L,\\
\hat{B}_2 &= B_{12} - A_{12}K^F.
\end{aligned}$$

Here we take the intraperiod leadership of the policymaker into account by substituting in constraint (22), but we treat u^L_t as given.

Because of the quadratic nature of the per-period objective in (18) and because policy and private-sector decisions are linear in the state, the discounted loss will necessarily have quadratic form in the state

$$V^F_t(\mathsf{y}_t) = \frac{1}{2}\mathsf{y}'_t S^F \mathsf{y}_t.$$

The Bellman equation characterizing discretionary policy of policymaker F, therefore, becomes

$$\begin{aligned}
\mathsf{y}'_t S^F \mathsf{y}_t = \min_{\mathsf{u}^F_t} \big(&\mathsf{y}'_t(\hat{Q} + \beta\hat{A}'S^F\hat{A})\mathsf{y}_t + 2\mathsf{y}'_t(\hat{P} + \beta\hat{A}'S^F\hat{B})\mathsf{u}_t\\
&+ \mathsf{u}'_t(\hat{R} + \beta\hat{B}'S^F\hat{B})\mathsf{u}_t\big).
\end{aligned} \tag{29}$$

Policy of the Leader. For the leader, policy u^L_t satisfies the following Bellman equation:

$$V^L_t(\mathsf{y}_t) = \min_{\mathsf{u}^L_t}\left(\frac{1}{2}\big(\mathsf{y}'_s\tilde{Q}\mathsf{y}_t + 2\mathsf{y}'_t\tilde{P}\mathsf{u}^L_t + \mathsf{u}^{L\prime}_t\tilde{R}\mathsf{u}^L_t\big) + \beta V^L_{t+1}\big(\tilde{A}\mathsf{y}_t + \tilde{B}\mathsf{u}^L_t\big)\right), \tag{30}$$

with

$$\begin{aligned}
\tilde{Q} = Q^L_{11} &- P^L_{12}G - G'P^{L\prime}_{12} + G'R^L_{22}F^F - Q^L_{12}\tilde{J} - \tilde{J}'Q^{L\prime}_{12}\\
&+ G'P^{L\prime}_{22}\tilde{J} + \tilde{J}'P^L_{22}G + \tilde{J}'Q^L_{22}\tilde{J},
\end{aligned}$$

$$\tilde{P} = P_{11}^L + \tilde{J}'Q_{22}^L\tilde{K} - Q_{12}^L\tilde{K} + G'P_{22}^{L\prime}\tilde{K} - P_{12}^L D - G'R_{21}^L$$
$$+ G'R_{22}^L D - \tilde{J}'P_{21}^L + \tilde{J}'P_{22}^L D,$$
$$\tilde{R} = R_{11}^L + \tilde{K}'Q_{22}^L\tilde{K} - R_{12}^L D - L'R_{21}^L + L'R_{22}^L D - \tilde{K}'P_{21}^L$$
$$+ \tilde{K}'P_{22}^L D - P_{21}^{L\prime}\tilde{K} + D'P_{22}^{L\prime}\tilde{K},$$
$$\tilde{J} = J - K^F G,$$
$$\tilde{K} = K^L - K^F D,$$
$$\tilde{A} = A_{11} - B_{12}G - A_{12}\tilde{J},$$
$$\tilde{B}_1 = B_{11} - B_{12}D - \tilde{A}_{12}\tilde{K},$$
$$\tilde{B}_2 = B_{21} - B_{22}D.$$

Here we take the intraperiod leadership of the policymaker into account by substituting in constraints (22) and (26). Similarly,

$$V_t^L(\mathsf{y}_t) = \frac{1}{2}\mathsf{y}_t'S^L\mathsf{y}_t,$$

and the Bellman equation characterizing discretionary policy of policymaker L becomes

$$\mathsf{y}_t'S^L\mathsf{y}_t = \min_{\mathsf{u}_t^L}\big(\mathsf{y}_s'(\tilde{Q} + \beta\tilde{A}'S^L\tilde{A})\mathsf{y}_t + 2\mathsf{y}_t'(\tilde{P} + \beta\tilde{A}'S^L\tilde{B})\mathsf{u}_t^L$$
$$+ \mathsf{u}_t^{L\prime}(\tilde{R} + \beta\tilde{B}'S^L\tilde{B})\mathsf{u}_t^L\big). \tag{31}$$

For a policy F^L, F^F and the private-sector response N, the evolution of the state variable satisfies the following equation:

$$\mathsf{y}_{t+1} = M\mathsf{y}_t, \tag{32}$$

where $M = A_{11} - A_{12}N - B_{11}F^L - B_{12}F^F$.

Discretionary Equilibrium as a Matrix Sextuple

Given y_0 and system matrices A and B, matrices N, F^L, G, and D define the trajectories $\{\mathsf{y}_s, \mathsf{x}_s, \mathsf{u}_s\}_{s=t}^{\infty}$ in a unique way and vice versa: if we know that $\{\mathsf{y}_s, \mathsf{x}_s, \mathsf{u}_s\}_{s=t}^{\infty}$ solve the discretionary optimization problem stated above, then, by construction, there are unique time-invariant linear relationships between them which we label as N, F^L, G, and D. Matrices S^L, S^F define the cost to go along a trajectory for each policymaker. Given the one-to-one mapping between equilibrium trajectories and $\{\mathsf{y}_s, \mathsf{x}_s, \mathsf{u}_s\}_{s=t}^{\infty}$ and the sextuple of matrices $\mathcal{T} = \{N, F^L, G, D, S^L, S^F\}$, it is convenient to continue with the definition of policy equilibrium in terms of $\mathcal{T}$, not trajectories.

The following proposition derives the first-order conditions for a discretionary optimization problem.

Proposition 3 (First-Order Conditions). *The first-order conditions to the optimization problem (17)–(18) can be written in the following form:*

$$\begin{aligned}
N &= (A_{22} + NA_{12})^{-1}((A_{21} - B_{22}F^F - B_{21}F^L \\
&\quad + N(A_{11} - B_{12}F^F - B_{11}F^L)), \qquad (33)\\
S^F &= \hat{Q} - 2\hat{P}_1 F^L - 2\hat{P}_2 F^F + F^{L\prime}\hat{R}_{11}F^L + 2F^{L\prime}\hat{R}_{12}F^F \\
&\quad + F^{F\prime}\hat{R}_{22}F^F + \beta\hat{A}'S^F\hat{A} - 2\beta\hat{A}'S^F(\hat{B}_1F^L + \hat{B}_2F^F) \\
&\quad + \beta F_1^{L\prime}\hat{B}'S^F(\hat{B}_1F^L + \hat{B}_2F^F) + \beta F^{F\prime}\hat{B}_2'S^F\hat{B}_2F^F, \qquad (34)\\
G &= \left(\hat{R}_{22} + \beta\hat{B}_2'S^F\hat{B}_2\right)^{-1}\left(\hat{P}_2' + \beta\hat{B}_2'S^F\hat{A}\right), \qquad (35)\\
D &= \left(\hat{R}_{22} + \beta\hat{B}_2'S^F\hat{B}_2\right)^{-1}\left(\hat{P}_1' + \beta\hat{B}_2'S^F\hat{B}_1\right), \qquad (36)\\
S^L &= \tilde{Q} + \beta\tilde{A}'S^L\tilde{A} - 2(\tilde{P} + \beta\tilde{A}'S^L\tilde{B})F^L \\
&\quad + F^{L\prime}(\tilde{R} + \beta\tilde{B}'S^L\tilde{B})F^L, \qquad (37)\\
F^L &= (\tilde{R} + \beta\tilde{B}'S^L\tilde{B})^{-1}(\tilde{P}' + \beta\tilde{A}'S^L\tilde{B}), \qquad (38)
\end{aligned}$$

where matrices F^F, $\hat{Q}$, $\hat{P}$, $\hat{R}$, $\hat{A}$, $\hat{B}$, $\tilde{Q}$, $\tilde{P}$, $\tilde{R}$, $\tilde{A}$, and $\tilde{B}$ are defined by (28) and (30).

Proof. From relationships (21), (22), and (27), it immediately follows that

$$N = J - K^F G + K^F D F^L - K^L F^L. \qquad (39)$$

A straightforward substitution of (23)–(25) and (27) into (39) leads to (33).

The optimal policy of policymaker F can be determined from (29) by differentiating the loss function with respect to u_t^F:

$$\mathsf{u}_t^F = -\big(\hat{R}_{22} + \beta\hat{B}_2'S^F\hat{B}_2\big)^{-1}\big(\big(\hat{P}_2' + \beta\hat{B}_2'S^F\hat{A}\big)\mathsf{y}_t + \big(\hat{P}_1' + \beta\hat{B}_2'S^F\hat{B}_1\big)\mathsf{u}_t^L\big),$$

from where the policymaker's reaction function is defined by (35) and (36). We simultaneously proved proposition 2. We substitute the reaction rules (19), (26), (35), and (36) in (29) and obtain equation (34).

The optimal policy (38) of policymaker L can be determined from (31) by differentiating the loss function with respect to u_t^L. We substitute the solution into (31) and obtain (37) for S^L.

DEFINITION 2. *The sextuple* $T = \{N, F^L, G, D, S^L, S^F\}$ *is a discretionary equilibrium under intraperiod leadership of policymaker L if it satisfies the system of first-order conditions (33)–(38).*

Definition 2 implicitly assumes that the first-order conditions are necessary and sufficient conditions of optimality. However, it is straightforward to demonstrate that under the assumption of symmetric positive semi-definite $\mathcal{Q}^i$, the second-order conditions for the minimum are almost surely satisfied for each policymaker; see, e.g., Blake and Kirsanova (2008).

Numerical Solution

One way to search for discretionary equilibrium is to use an algorithm similar to the one in Oudiz and Sachs (1985) and Backus and Driffill (1986). We initialize the matrices $\{N_0, S_0^L, S_0^F\}$ and then solve the non-linear system of first-order conditions (33)–(38) using an appropriate iterative scheme. This algorithm can be interpreted as search for equilibria that are "iterative-expectations stable" under joint learning; see Dennis and Kirsanova (2010). Alternatively, one can iterate between the solution to (33) that describes the response of the private-sector given policy, and the solution to (34)–(38) that describes the best policy given the response of the private sector; see Blake and Kirsanova (2008).

Appendix 2. Social Welfare

The procedure to derive the appropriate welfare metric is standard and explained in Woodford (2003a). The one-period (flow) welfare in (14) is $\mathcal{W}_t$:

$$\mathcal{W}_t = u(C_t) + f(G_t) - \int_0^1 v(y_t(z))dz.$$

Around the steady state, a quadratic approximation to this is

$$\mathcal{W}_t = Cu_C(C)\left(\hat{C}_t + \frac{1}{2}\left(1 - \frac{1}{\sigma}\right)\hat{C}_t^2\right) + Gf_G(G)\left(\hat{G}_t + \frac{1}{2}\left(1 - \frac{1}{\sigma}\right)\hat{G}_t^2\right)$$
$$- Y v_y(Y)\left(\hat{Y}_t + \frac{1}{2}\left(1 + \frac{1}{\psi}\right)\hat{Y}_t^2 + \frac{1}{2}\left(\frac{1}{\psi} + \frac{1}{\epsilon}\right) var_z \hat{y}_t(z)\right). \tag{40}$$

Further, a second-order approximation of aggregate demand (8) can be written as

$$\hat{C} = \frac{1}{\theta}\left(\hat{Y} - (1-\theta)\hat{G} - \theta\frac{1}{2}\hat{C}^2 - \frac{1}{2}(1-\theta)\hat{G}^2 + \frac{1}{2}\hat{Y}^2\right),$$

so we can substitute consumption in (40) and obtain

$$\mathcal{W}_s = \theta u_C\left(\left(1 - \frac{v_y}{u_C}\right)\hat{Y}_s - (1-\theta)\left(1 - \frac{f_G}{u_C}\right)\hat{G}_s\right.$$
$$- \frac{(1-\theta)}{2}\left(1 + \frac{f_G}{u_C}\frac{(1-\sigma)}{\sigma}\right)\hat{G}_s^2 - \frac{\theta}{2\sigma}\hat{C}^2$$
$$\left.-\frac{1}{2}\left(\frac{v_y}{u_C}\frac{1+\psi}{\psi} - 1\right)\hat{Y}_s^2 - \frac{1}{2}\frac{v_y}{u_C}\frac{\psi+\epsilon}{\psi\epsilon} var_z \hat{y}_s(z)\right).$$

To transform this equation into a more convenient form that does not include linear terms, we proceed as follows (see Beetsma and Jensen 2005). If the government removes monopolistic distortions *and* distortions from income taxation in the steady state using a subsidy[20] $\mu^w = \frac{\mu}{(1-\tau)}$, then $u_C/v_y = f_G/u_C = 1$ so the linear terms in (40) drop out. The final formula for social welfare is

$$\mathcal{W}_s = -\theta u_C \left(\frac{\theta}{2\sigma} c_s^2 + \frac{(1-\theta)}{2\sigma} g_s^2 + \frac{1}{2\psi} y_s^2 + \frac{1}{2} \left(\frac{1}{\psi} + \frac{1}{\epsilon} \right) var_z \hat{y}_s(z) \right).$$

Woodford (2003a) has shown that

$$\sum_{t=0}^{\infty} \beta^t var_z \hat{y}_s(z) = \sum_{t=0}^{\infty} \beta^t \frac{\gamma \epsilon^2}{(1-\gamma\beta)(1-\gamma)} \pi_t^2,$$

so, using a conventional notation for gap variables, we get the final formula for the social welfare function:

$$\begin{aligned} \mathcal{W}_s &= -\frac{\theta(\epsilon+\psi)\gamma\epsilon u_C}{2\psi(1-\gamma\beta)(1-\gamma)} \left(\frac{\kappa}{\epsilon} \left(\frac{\theta}{\sigma} c_t^2 + \frac{(1-\theta)}{\sigma} g_t^2 + \frac{1}{\psi} y_t^2 \right) + \pi_t^2 \right) \\ &= -\frac{\theta(\epsilon+\psi)\gamma\epsilon u_C}{2\psi(1-\gamma\beta)(1-\gamma)} \left(\rho_c c_s^2 + \rho_g g_s^2 + \rho_y y_s^2 + \pi_s^2 \right), \end{aligned}$$

which, after normalization, is (15) in the main text with $\rho_c = \frac{\kappa\theta}{\epsilon\sigma}$, $\rho_g = \frac{\kappa(1-\theta)}{\epsilon\sigma}$, $\rho_y = \frac{\kappa}{\psi\epsilon}$.

References

Adam, K., and R. M. Billi. 2008. "Monetary Conservatism and Fiscal Policy." *Journal of Monetary Economics* 55 (8): 1376–88.

Anderson, E. W., L. P. Hansen, E. R. McGrattan, and T. J. Sargent. 1996. "Mechanics of Forming and Estimating Dynamic Linear Economies." In *Handbook of Computational Economics*, ed. H. M. Amman, D. A. Kendrick, and J. Rust, 171–252. Elsevier.

Auerbach, A. J. 2003. "Is There a Role for Discretionary Policy?" In *Rethinking Stabilization Policy*. Federal Reserve Bank of Kansas City.

Backus, D., and J. Driffill. 1986. "The Consistency of Optimal Policy in Stochastic Rational Expectations Models." CEPR Discussion Paper No. 124.

[20]This subsidy is financed, of course, by lump-sum taxation, $\mu^w = T$.

Barro, R., and D. Gordon. 1983. "Rules, Discretion and Reputation in a Model of Monetary Policy." *Journal of Monetary Economics* 12 (1): 101–21.

Beetsma, R. M., and A. L. Bovenberg. 2006. "Political Shocks and Public Debt: The Case for a Conservative Central Bank Revisited." *Journal of Economic Dynamics and Control* 30 (11): 1857–83.

Beetsma, R., and H. Jensen. 2004. "Mark-Up Fluctuations and Fiscal Policy Stabilization in a Monetary Union." *Journal of Macroeconomics* 26 (2): 357–76.

———. 2005. "Monetary and Fiscal Policy Interactions in a Micro-Founded Model of a Monetary Union." *Journal of International Economics* 67 (2): 320–52.

Benigno, P., and M. Woodford. 2004. "Optimal Monetary and Fiscal Policy: A Linear-Quadratic Approach. In *NBER Macroeconomics Annual 2003*, ed. M. Gertler and K. Rogoff, 271–333. Cambridge, MA: MIT Press.

Blake, A. P., and T. Kirsanova. 2008. "Discretionary Policy and Multiple Equilibria in LQ RE Models." Mimeo, University of Exeter. Available at SSRN, http://ssrn.com/abstract=943032.

Canzoneri, M., R. Cumby, B. Diba, and O. Mykhaylova. 2006. "New Keynesian Explanations of Cyclical Movements in Aggregate Inflation and Regional Inflation Differentials." *Open Economies Review* 17 (1): 27–55.

Clarida, R. H., J. Galí, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661–1707.

Cooper, R., and A. John. 1988. "Coordinating Coordination Failures in Keynesian Models." *Quarterly Journal of Economics* 103 (3): 441–63.

Currie, D., and P. Levine. 1993. *Rules, Reputation and Macroeconomic Policy Coordination*. Cambridge: Cambridge University Press.

de Zeeuw, A. J., and F. van der Ploeg. 1991. "Difference Games and Policy Evaluation: A Conceptual Framework." *Oxford Economic Papers* 43 (4): 612–36.

Dennis, R., and T. Kirsanova. 2010. "Expectations Traps and Coordination Failures: Selecting Among Multiple Discretionary Equilibria." Federal Reserve Bank of San Francisco Working Paper No. 2.

Dixit, A., and L. Lambertini. 2003. "Interactions of Commitment and Discretion in Monetary and Fiscal Policies." *American Economic Review* 93 (5): 1522–42.

Dixit, A., and J. Stiglitz. 1977. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review* 67 (3): 297–308.

Favero, C., and T. Monacelli. 2005. "Fiscal Policy Rules and Regime (In)Stability: Evidence from the U.S." Mimeo, University of Bocconi.

Hughes Hallett, A., J. Libich, and P. Stehlik. 2009. "Rogoff Revisited: The Conservative Central Banker Proposition under Active Fiscal Policies." *Economic Letters* 104 (3): 140–43.

Ireland, P. N. 2004. "Technology Shocks in the New Keynesian Model." *Review of Economics and Statistics* 86 (4): 923–36.

Lambertini, L. 2006. "Monetary-Fiscal Interactions with a Conservative Central Bank." *Scottish Journal of Political Economy* 53 (1): 90–128.

Leeper, E. M. 1991. "Equilibria Under 'Active' and 'Passive' Monetary and Fiscal Policies." *Journal of Monetary Economics* 27 (1): 129–47.

Oudiz, G., and J. Sachs. 1985. "International Policy Coordination in Dynamic Macroeconomic Models." In *International Economic Policy Coordination*, ed. W. H. Buiter and R. C. Marston, 274–318. Cambridge: Cambridge University Press.

Rogoff, K. 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics* 100 (4): 1169–89.

Rudebusch, G. 2002. "Assessing Nominal Income Rules for Monetary Policy with Model and Data Uncertainty." *Economic Journal* 112 (479): 402–32.

Smets, F., and R. Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1 (5): 1123–75.

Sutherland, A. 2002. "A Simple Second-Order Solution Method for Dynamic General Equilibrium Models." CEPR Discussion Paper No. 3554.

Svensson, L. E. O. 1997. "Optimal Inflation Targets, 'Conservative' Central Banks, and Linear Inflation Contracts." *American Economic Review* 87 (1): 98–114.

Taylor, J. B. 2000. "Reassessing Discretionary Fiscal Policy." *Journal of Economic Perspectives* 14 (3): 21–36.

Vestin, D. 2006. "Price-Level versus Inflation Targeting." *Journal of Monetary Economics* 53 (7): 1361–76.

Walsh, C. 2003. "Speed Limit Policies: The Output Gap and Optimal Monetary Policy." *American Economic Review* 93 (1): 265–78.

Woodford, M. 2003a. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.

———. 2003b. "Optimal Interest-Rate Smoothing." *Review of Economic Studies* 70 (4): 861–86.

Endogenous Central Bank Information and the Optimal Degree of Transparency*

Romain Baeriswyl
Munich Graduate School of Economics and Swiss National Bank

As a policymaker, the central bank both observes and shapes the economy. The central bank scrutinizes market activity to assess the state of the economy, and its policy strongly shapes market outcomes. When transparency allows the central bank to shape the economy more effectively, it may also cause the informational role of the economic aggregate to deteriorate. This paper presents a simple model to capture the endogenous nature of central bank information and to address welfare issues. First, accounting for the endogeneity of information highlights the detrimental effects of transparency. A model with endogenous information always calls for a lower degree of transparency than a model with exogenous information. Second, the optimal degree of transparency for endogenous information is unrelated to the accuracy of firms' private information.

JEL Codes: D82, E52, E58.

1. Introduction

As a policymaker, the central bank both observes and influences the economy. On the one hand, the central bank observes market activity to assess the state of the economy and to decide the course of its policy. In practice, a central bank devotes extensive resources to collecting data on economic agents' behavior to estimate aggregate economic outcomes. For example, to make its policy decisions, a

*The author thanks Camille Cornand, Frank Heinemann, Gerhard Illing, an anonymous referee, and Carl Walsh, the editor, for very helpful discussions and comments. Remaining errors are the author's own. The views expressed in this paper do not necessarily reflect those of the Swiss National Bank. The author gratefully acknowledges financial support from the German Research Foundation (DFG) through GRK 801. Author e-mail: Romain.Baeriswyl@snb.ch.

central bank will predict price levels to estimate potential economic imbalances. In a prominent article, Hayek (1945) emphasized the informational role of prices and argued that prices are not only an exchange rate between goods but also an information aggregator. He pointed out that the prices determined by decentralized markets are an essential source of information because they aggregate the "dispersed bits of incomplete and frequently contradictory knowledge, which all the separate individuals possess."

On the other hand, the policy that is implemented by the central bank influences the economy, particularly private-sector expectations. Indeed, the shaping of market expectations plays a key role in the conduct of monetary policy. Woodford (2005) pointed out, "for [monetary policy to be most effective] not only do expectations about policy matter, but [...] very little else matters."

However, the ambivalent role of the central bank—as an observer and as a shaper—causes a dilemma concerning the implementation of monetary policy, which has been documented by Amato and Shin (2006) and Morris and Shin (2005); the more successfully a central bank influences market expectations, the less reliably market outcomes serve as indicators of the state of the economy. Although market outcomes would reflect the true state of the economy in the absence of central bank interventions (i.e., if the central bank were an observer only), they nevertheless partly reflect the central bank's expectation because it intervenes in economic development. As soon as the central bank's estimation errors cause it to misread the state of the economy, these errors will be reflected in market outcomes. This dilemma raises the question of the desirability of central bank transparency.

Although the current paper discusses the theoretical possibility that central bank transparency spoils the informational content of economic aggregates, some empirical analyses highlight the plausibility of the mechanism.[1] For instance, in their analysis of U.S. data, Ehrmann and Fratzscher (2005) showed that with increasing transparency "markets attach more importance to the statements and the balance-of-risk assessments at FOMC meetings and less importance to news about macroeconomic fundamentals." They conclude that

[1] See Geithner (2006) for a general discussion on central bank uncertainty and transparency.

"the reaction of financial markets to the release of macroeconomic fundamentals can be an important source of information for the central bank about the markets' diverse and possibly deviating views," and that "under its new disclosure policy, the Federal Reserve has less such information available." Transparency standards in central banking have improved in recent decades, as Geraats (2009) documented, and therefore this argument should have become more relevant.

This paper aims to develop a simple model with strategic complementarities under imperfect common knowledge that captures the endogenous nature of central bank information. Central bank information is endogenous in the sense that its accuracy is a function of the disclosure strategy of the central bank itself.[2] In this context, increasing the level of ambiguity in the central bank's disclosure also increases the accuracy of its information. Moreover, increasing the level of ambiguity in the central bank's disclosure may even increase the accuracy of firms' information.

The feedback effect between the central bank and the private sector, which arises from endogenous information, is reminiscent of the case made by Bernanke and Woodford (1997). They addressed the existence and uniqueness of a rational expectations equilibrium when the central bank observes and responds to private-sector forecasts. In particular, they showed that a central bank cannot infer the value of the state variable by observing private-sector forecasts and simultaneously stabilize the economy fully. This situation arises because if inflation equals the target in equilibrium, then the information of the private forecasters is not revealed. It was the stabilizing action of the central bank that rendered the private-sector forecasts uninformative in Bernanke and Woodford's paper, but, in the present analysis, it is the noisy announcement of the central bank that makes the aggregated price level less informative.

[2]The endogeneity of central bank information is to be distinguished from the endogeneity of central bank release. The information released by the action of the central bank to the private sector is endogenous in the sense that it depends on the central bank's preferences and economic assessment. However, this does not imply that the information of the central bank is endogenous. For instance, Walsh (2007) or Baeriswyl and Cornand (2010) analyze the conduct of monetary policy when the central bank release is endogenous but its information is exogenous.

From a normative point of view, this paper contributes to the ongoing debate on the welfare effects of central bank transparency and emphasizes that the deterioration of central bank information (and disclosure) is a potentially detrimental effect. This argument resonates with that of Morris and Shin (2002) (hereafter M-S). In their seminal beauty-contest paper on exogenous information, they highlighted the potentially detrimental effect of noisy public information. In an environment of strategic complementarities, central bank disclosure is given too much weight relative to its face value because it serves as a focal point. Higher-order expectations are mainly driven by public disclosures, and therefore the response to a noisy public disclosure, exacerbated by the coordination motive, may destabilize the economy. Thus, reducing the degree of transparency improves welfare because it reduces the extent of common knowledge about the disclosure and the weight assigned to it for higher-order expectations.

Moreover, endogenous information addresses the detrimental effect of transparency to the accuracy of central bank disclosure and thereby of firms' information. M-S focuses on the large weight assigned to public disclosures for higher-order expectations, which are exacerbated by coordination, but the present model of endogenous information emphasizes the negative accuracy effect of transparency for both first- and higher-order expectations.

Two conclusions can be drawn with respect to the optimal degree of transparency. First, accounting for the endogeneity of information reinforces the detrimental effect of transparency. A model with endogenous information always calls for a lower degree of transparency than a model with exogenous information. Second, the optimal degree of transparency for endogenous information is unrelated to the accuracy of firms' private information. This conclusion contrasts with M-S's model for exogenous information, in which a decrease in the accuracy of firms' private information increased the optimal degree of transparency.

Section 2 describes the economy such that the Keynesian beauty contest, formalized by M-S, is interpreted as the price-setting problem of monopolistically competitive firms.[3] Section 3 presents the

[3]See also Amato, Morris, and Shin (2002), Hellwig (2005), or Hellwig and Veldkamp (2009) for an application of M-S to price-setting decisions by firms.

model for exogenous central bank information and replicates M-S's analysis as a benchmark. Section 4 derives the model for endogenous information and discusses the effects of transparency on central bank information and disclosure. While sections 3 and 4 address the welfare effects of transparency on a broad range of welfare functions, section 5 focuses on microfounded welfare, which highly weights coordination at the social level. The conclusions are presented in section 6.

2. The Economy

The economy is populated by a representative household, a *continuum* of monopolistic competitive firms, and a central bank. We abstract from the microfounded market interactions since they are very standard and focus on the optimal behavior of firms.[4]

2.1 Firms

In an economy where the representative household consumes a composite good à la Dixit-Stiglitz and where goods are imperfect substitutes, the optimal pricing rule of firm i is given by

$$p_i = \mathbb{E}_i[p + \xi c], \tag{1}$$

where $\mathbb{E}_i$ is the expectation operator of firm i conditional on its information, p is the overall price level, and c is the real output gap. The pricing rule (1) says that each firm sets its price according to both its own belief about the real output gap and its belief about the overall price level. We assume that the nominal aggregate demand defined as $c + p$ is determined by a stochastic demand shock $g \in \mathbb{R}$. So, one can write the pricing rule as

$$p_i = \mathbb{E}_i[(1 - \xi)p + \xi g]. \tag{2}$$

The parameter ξ captures the impact of the real output gap on prices (through wages). A large ξ means that the representative household is highly risk averse and that output gaps imply large variations in

[4]See Adam (2007) for a full derivation of the microfoundations.

wages and therefore in prices. We shall assume in this paper that $0 < \xi < 1$, which implies that prices are strategic complements, meaning that firms tend to raise their price whenever they expect the others to do so. This assumption seems very natural and captures the concept of beauty contest introduced by Keynes: firms base their decision not only on their own expectations of fundamentals but also on the so-called higher-order expectations, i.e., on their expectation of others' expectations of fundamentals.

Substituting successively the average price level with higher-order expectations about the demand shock, the pricing rule becomes

$$\begin{aligned} p_i &= \mathbb{E}_i[(1-\xi)p + \xi g] \\ &= \mathbb{E}_i\Big[\xi g + (1-\xi)\left[\bar{\mathbb{E}}[\xi g + (1-\xi)[\bar{\mathbb{E}}[\xi g + \ldots]]]\right]\Big]. \end{aligned} \tag{3}$$

Using the fact that with heterogeneous information the law of iterated expectations fails since expectations of higher order do not collapse to the average expectation of degree one,[5] the pricing rule can be rewritten as

$$p_i = \xi \sum_{k=0}^{\infty} (1-\xi)^k \mathbb{E}_i[\bar{\mathbb{E}}^{(k)}(g)],$$

and averaging over firms, we get

$$p = \xi \sum_{k=0}^{\infty} (1-\xi)^k [\bar{\mathbb{E}}^{(k+1)}(g)], \tag{4}$$

where k is the degree of higher-order iteration, and $\bar{\mathbb{E}}$ is the population average expectation operator such that $\bar{\mathbb{E}}(.) = \int_i \mathbb{E}_i(.)di$. We use the following notation of higher-order expectations: $\bar{\mathbb{E}}^{(0)}(x) = x$ is the expected variable x itself, $\bar{\mathbb{E}}^{(1)}(x) = \bar{\mathbb{E}}(x)$ is the average expectation of x, $\bar{\mathbb{E}}^{(2)}(x) = \bar{\mathbb{E}}\bar{\mathbb{E}}^{(1)}(x) = \bar{\mathbb{E}}\bar{\mathbb{E}}(x)$ is the average expectation of the average expectation of x, and so on.

To take its pricing decision, each firm receives two signals. First, each firm gets a private signal about the demand shock. The private

[5] See Morris and Shin (2002).

signal is centered on the true value of g and has a normally distributed error term:

$$g_i = g + \varepsilon_i \quad \text{with } \varepsilon_i \sim N(0, \sigma_\varepsilon^2), \tag{5}$$

where ε_i are identically and independently distributed across firms.

Second, the central bank provides firms with its viewpoint about the demand shock. The central bank communicates its information D with more or less ambiguity. We capture this ambiguity with the degree of transparency of its disclosure. Generally speaking, transparency refers to the case where the central bank shares its information with the private sector (symmetric information).[6] Imperfect transparency is commonly interpreted in terms of noise variance that induces uncertainty to the private sector on the central bank's assessment.[7] With heterogenous information, transparency is interpreted as the degree of common knowledge among private agents. Heinemann and Illing (2002) propose to control the degree of common knowledge by introducing idiosyncratic noise in central bank disclosure.

The signal disclosed by the central bank and received by firm i is written as

$$D_i = D + \phi_i \quad \text{with } \phi_i \sim N(0, \sigma_\phi^2). \tag{6}$$

The dispersion of individual noises σ_ϕ^2 determines the degree of transparency of the central bank. Under transparency, every firm gets the same univocal signal ($\sigma_\phi^2 = 0$). Then, the central bank disclosure D is a public signal that is common knowledge among firms. Under opacity, the individual signal gotten by each firm has an infinite idiosyncratic noise ($\sigma_\phi^2 \to \infty$). The central bank disclosure thus does not contain any valuable information.

---

[6]According to the classification of Geraats (2002), the transparency discussed in this paper is economic transparency since the only uncertainty in the economy is the value of the fundamental shock g.

[7]For instance, Cukierman and Meltzer (1986) interpret transparency as the noise variance of monetary control or Faust and Svensson (2002) as the noise variance of information on monetary control errors. See Geraats (2002) for an overview.

2.2 Welfare

In an economy characterized by monopolistic competition, the welfare of the representative household is decreasing in both the dispersion of prices across firms $\int_i (p_i - p)^2 di$ and the variability of the output gap $c = g - p$. Therefore, we define the social loss as

$$L = \int_i (p_i - p)^2 di + \lambda (g - p)^2, \tag{7}$$

where λ is the weight assigned to the output-gap variability. As Angeletos and Pavan (2007) stressed, demand shocks create inefficiencies only when information is incomplete.[8]

The welfare function used in the transparency debate of M-S is a controversial matter because the detrimental effect of transparency is driven by the relative relevance of coordination (dispersion) and stabilization (distortion) at the social level. However, the application of the M-S argument to different welfare functions may lead to different conclusions. For example, Hellwig (2005) and Woodford (2005) showed that when coordination is socially highly valuable, transparency improves welfare because it helps to coordinate firms' price setting. We leave λ unspecified in sections 3 and 4 to discuss the M-S argument in an environment where coordination is not particularly socially valuable and to emphasize the effect of endogenous information in this context. We show in section 3.2 that the welfare in M-S given by $-\int_i (p_i - g)^2 di$ can be expressed by (7) with the parameter $\lambda = 1$. Thus, the welfare function in M-S gives equal weight to coordination and stabilization at the social level. Then, in section 5, following Hellwig (2005) and Woodford (2005), we consider the welfare function that is consistent with a microfounded economy.

2.3 The Central Bank

The central bank seeks to minimize the unconditional expected loss (7) by disclosing information to firms about the fundamental

[8]Indeed, with perfect information, $p_i = p = g$ ensures that the loss (7) is zero.

demand shock g.[9] We will discuss the welfare effect of the central bank's disclosure in two informational cases.

First, we consider the case where the central bank directly observes the stochastic demand shock g with some noise. The precision of central bank information is then exogenously determined and independent of its disclosure strategy (section 3). Second, we assume that the central bank cannot directly observe the demand shock g but instead watches market activity to evaluate the state of the economy. In this case, we show that the precision of central bank information is endogenous because it depends upon its disclosure strategy (section 4).

For the sake of generality, we define the central bank's information as D and the variance of the central bank's expectation error about the fundamental shock g as

$$\mathrm{Var}[\mathbb{E}(g|D) - g] \equiv \sigma_\mu^2. \tag{8}$$

This definition allows us to solve generally for the equilibrium behavior of firms before specifying whether the central bank's information is exogenous or endogenous.

2.4 Equilibrium

This section derives the perfect Bayesian equilibrium behavior of firms. To determine the optimal price rule (4), we build the first- and higher-order expectations of firm i about the demand shock g conditional on its information. Given firms' information (5), (6), and (8), the expectation of degree one about the demand shock $\mathbb{E}_i(g)$ yields

$$\mathbb{E}(g|g_i, D_i) = \frac{\sigma_\mu^2 + \sigma_\phi^2}{\sigma_\varepsilon^2 + \sigma_\mu^2 + \sigma_\phi^2} g_i + \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\mu^2 + \sigma_\phi^2} D_i = \Omega_{11} g_i + \Omega_{12} D_i. \tag{9}$$

[9] Note that the central bank does not implement any monetary instrument to stabilize the economy because the instrument would be indeterminate since the price level does not enter the loss function as defined in section 2.2. See Baeriswyl and Cornand (2010) for an analysis that accounts for the interaction between the monetary instrument and the disclosure of the central bank.

The best estimate of the demand shock by firm i is an average of both its signals whose weighting depends upon their relative precision. To compute the higher-order expectations of firm i, one needs also to know the expectation of degree one of the central bank average disclosure $\mathbb{E}_i(D)$. This delivers

$$\mathbb{E}(D|g_i, D_i) = \frac{\sigma_\phi^2}{\sigma_\varepsilon^2 + \sigma_\mu^2 + \sigma_\phi^2} g_i + \frac{\sigma_\varepsilon^2 + \sigma_\mu^2}{\sigma_\varepsilon^2 + \sigma_\mu^2 + \sigma_\phi^2} D_i = \Omega_{21} g_i + \Omega_{22} D_i. \tag{10}$$

Note that under transparency (when $\sigma_\phi^2 = 0$), the central bank's disclosure is univocal and $\Omega_{21} = 0$, which means that the private signal g_i does not help in guessing D (since $D_i = D$). Under opacity, when the idiosyncratic noise is infinite ($\sigma_\phi^2 \to \infty$), the central bank's disclosure is of no use to estimate the demand shock g and the best estimate is the private signal g_i itself ($\Omega_{11} = 1$).

Using these results, we can express the higher-order expectation of degree k as

$$\bar{\mathbb{E}}^{(k)} \begin{pmatrix} g \\ D \end{pmatrix} = \begin{pmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{21} & \Omega_{22} \end{pmatrix}^k \begin{pmatrix} g \\ D \end{pmatrix}.$$

Plugging this into the price rule (4), we get

$$p = \begin{pmatrix} \xi & 0 \end{pmatrix} \sum_{k=0}^{\infty} (1-\xi)^k \begin{pmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{21} & \Omega_{22} \end{pmatrix}^{k+1} \begin{pmatrix} g \\ D \end{pmatrix}. \tag{11}$$

The price rule is a linear combination of the demand shock and the central bank's average disclosure:

$$\begin{aligned} p &= \gamma_1 g + \gamma_2 D \quad \text{with} \\ \gamma_1 &= \frac{(1-\xi)\Omega_{21}\gamma_2 + \xi\Omega_{11}}{1-(1-\xi)\Omega_{11}} = \frac{\xi\sigma_\mu^2 + \sigma_\phi^2}{\sigma_\varepsilon^2 + \xi\sigma_\mu^2 + \sigma_\phi^2} \\ \gamma_2 &= \frac{(1-\xi)\Omega_{12}\gamma_1 + \xi\Omega_{12}}{1-(1-\xi)\Omega_{22}} = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \xi\sigma_\mu^2 + \sigma_\phi^2}. \end{aligned} \tag{12}$$

γ_1 and γ_2 sum up to 1. The equilibrium firms' action can be interpreted as a weighted average of the fundamental g and the average disclosure D. The weight assigned to the central bank's disclosure is larger in the equilibrium action (12) than in the best estimate of g given in (9): $\gamma_2 > \Omega_{12}$. This discrepancy arises because of the coordination motive in the pricing rule. While ε_i and ϕ_i are idiosyncratic noises, the central bank noise with variance σ_μ^2 is commonly observed by all firms through the disclosure D_i. The weight assigned to the central bank's error (and thereby to D_i) increases as the coordination motive strengthens; strategic complementarities raise the firms' incentives to coordinate their actions around the central bank's disclosure. When the degree of strategic complementarities $1-\xi$ increases, the weight assigned to the private signal g_i declines ($\frac{\partial\gamma_1}{\partial\xi} > 0$), and the weight assigned to the central bank's disclosure increases ($\frac{\partial\gamma_2}{\partial\xi} < 0$). When the degree of transparency increases (σ_ϕ^2 falls), the weight given to the central bank's disclosure D_i increases because firms can interpret it less ambiguously and better guess the actions of others ($\frac{\partial\gamma_1}{\partial\sigma_\phi^2} > 0$ and $\frac{\partial\gamma_2}{\partial\sigma_\phi^2} < 0$). Signals are also given a higher weight when their precision increases: $\frac{\partial\gamma_1}{\partial\sigma_\varepsilon^2} < 0$ and $\frac{\partial\gamma_2}{\partial\sigma_\mu^2} < 0$.

3. Exogenous Central Bank Information

This section analyzes the welfare effect of central bank disclosure when the central bank directly observes the demand shock. The aim of this section is to illustrate M-S's much-debated conclusion where central bank information is exogenous. The present section should be seen as a benchmark case that replicates M-S's results and allows a better comparison with the endogenous case in section 4.

We describe the information structure and discuss the optimal information disclosure for two cases. We first examine the case when the central bank chooses between full transparency and full opacity (i.e., the central bank either perfectly reveals its opinion or totally withholds it) and, second, the case when the central bank can choose its optimal degree of transparency (i.e., the central bank makes announcements with some ambiguity).

3.1 Information Structure

Under exogenous information, the central bank directly (but imperfectly) observes the demand shock g. According to the definition of the error term of central bank information (8), we assume that the central bank receives a signal D on the demand shock that is centered on its true value g and contains an error term μ:

$$D = g + \mu \quad \text{with } \mu \sim N(0, \sigma_\mu^2).$$

The precision of central bank information σ_μ^2 is exogenous.

3.2 Welfare

Given the equilibrium behavior of firms (12) and the central bank information as described in the previous section, the unconditional expected social loss (7) can be written as

$$\mathbb{E}(L) = \gamma_1^2\sigma_\varepsilon^2 + \gamma_2^2\sigma_\phi^2 + \lambda\gamma_2^2\sigma_\mu^2 = \frac{\sigma_\varepsilon^2(\lambda\sigma_\mu^2 + \sigma_\phi^2) + (\xi\sigma_\mu^2 + \sigma_\phi^2)^2}{(\sigma_\varepsilon^2 + \xi\sigma_\mu^2 + \sigma_\phi^2)^2}\sigma_\varepsilon^2. \tag{13}$$

The welfare considered in M-S is given by $-\int_i (p_i - g)^2 di$. We write the corresponding loss as

$$\mathbb{E}(L_{MS}) = \mathbb{E}\left(\int_i (p_i - g)^2 di\right) = \gamma_1^2\sigma_\varepsilon^2 + \gamma_2^2\sigma_\phi^2 + \gamma_2^2\sigma_\mu^2.$$

This implies that the welfare in M-S is a particular case of our general formulation (13) where $\lambda = 1$. This means that the model of M-S equally weights coordination and stabilization at the social level.

3.3 Transparency versus Opacity

3.3.1 Opacity

The welfare is now computed when the central bank withholds its information, i.e., $\sigma_\phi^2 \to \infty$. Under opacity, firms set their price equal to their private signal g_i, i.e., $\gamma_1 = 1$ and $\gamma_2 = 0$. The resulting expected loss is

$$\mathbb{E}(L_O) = \mathbb{E}\left(\int_i (\gamma_1(g + \varepsilon_i) - \gamma_1 g)^2 di + \lambda(g - \gamma_1 g)^2\right) = \sigma_\varepsilon^2.$$

The overall price level p is equal to the fundamental g, which implies an output gap of zero. The price dispersion across firms is given by the variance of the idiosyncratic noise ε_i.

3.3.2 Transparency

Under transparency, the disclosure of the central bank is common knowledge ($\sigma_\phi^2 = 0$) and the pricing rule of firms becomes

$$p = \frac{\xi\sigma_\mu^2}{\sigma_\varepsilon^2 + \xi\sigma_\mu^2} g + \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \xi\sigma_\mu^2} D,$$

which yields the expected loss

$$\mathbb{E}(L_T) = \left(\frac{\xi\sigma_\mu^2}{\sigma_\varepsilon^2 + \xi\sigma_\mu^2}\right)^2 \sigma_\varepsilon^2 + \lambda\left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \xi\sigma_\mu^2}\right)^2 \sigma_\mu^2.$$

Transparency improves welfare when the loss under opacity L_O is larger than the loss under transparency L_T. When the precision of central bank information is exogenous, full transparency is preferable to opacity when

$$\lambda - 2\xi < \psi, \tag{14}$$

where $\psi = \frac{\sigma_\varepsilon^2}{\sigma_\mu^2}$ is the relative inaccuracy of private and central bank information. Transparency is detrimental to welfare when public information is too noisy relative to private information, when the degree of strategic complementarities is relatively high, and when the weight assigned to coordination is relatively low.

The general framework developed in this paper shows the extent to which the welfare effect of transparency is related to the social value of coordination. In the case of M-S, as $\lambda = 1$, private information must be more accurate than public information ($\psi < 1$) for transparency to be detrimental. Geraats (2002) and Svensson (2006) argue that the detrimental effect of transparency emphasized in M-S's beauty-contest framework arises under unrealistic conditions because the information held by public institutions (e.g., a central bank) is typically more accurate than the information that is privately available.[10] However, if the social value of coordination

[10]For instance, in an empirical analysis on U.S. data, Romer and Romer (2000) show that the Federal Reserve better forecasts the output and inflation than any single private commercial bank.

is smaller than in M-S ($\lambda > 1$), opacity may be preferable even when public information is more accurate than private information ($\lambda - 2\xi > 1$).

3.4 Optimal Degree of Transparency

In the previous section, the central bank could either disclose its noisy information with perfect precision or withhold it. In reality, however, central bankers are known for their ambiguous mumblings. Central bank disclosures are, therefore, open to interpretation. The more equivocally a central bank discloses information, the higher the uncertainty surrounding both the interpretation of the disclosure (fundamental uncertainty) and its interpretation by others (strategic uncertainty). When full transparency is detrimental to welfare relative to opacity, reducing transparency may improve welfare. However, even when full transparency is preferable to opacity, partial transparency may still yield a superior outcome.[11]

To determine the optimal degree of transparency σ_ϕ^2, we minimize the loss (13) with respect to σ_ϕ^2 and set it equal to zero:

$$\begin{aligned}\frac{\partial \mathbb{E}(L)}{\partial \sigma_\phi^2} &= 2\gamma_1\sigma_\varepsilon^2\frac{\partial\gamma_1}{\partial\sigma_\phi^2} + \gamma_2^2 + 2\gamma_2\sigma_\phi^2\frac{\partial\gamma_2}{\partial\sigma_\phi^2} + 2\lambda\gamma_2\sigma_\mu^2\frac{\partial\gamma_2}{\partial\sigma_\phi^2} \\ &= \frac{(\sigma_\varepsilon^2 + (3\xi - 2\lambda)\sigma_\mu^2 + \sigma_\phi^2)\sigma_\varepsilon^4}{(\sigma_\varepsilon^2 + \xi\sigma_\mu^2 + \sigma_\phi^2)^3} \\ &= 0 \quad \Leftrightarrow \quad \sigma_\phi^2 = (2\lambda - 3\xi)\sigma_\mu^2 - \sigma_\varepsilon^2. \end{aligned} \tag{15}$$

Since the variance of idiosyncratic noise is non-negative, the optimal degree of transparency is described by

$$\chi_{\text{exo}} = \max[0; 2\lambda - 3\xi - \psi], \tag{16}$$

where $\chi_{\text{exo}} = \frac{\sigma_\phi^2}{\sigma_\mu^2}$ is the relative variance of idiosyncratic and fundamental noise of the central bank's disclosure. This analysis calls for

[11] This idea goes back to Cornand and Heinemann (2008), who show that reducing the degree of common knowledge is welfare improving. However, they propose to reduce the degree of common knowledge by disclosing a public signal to a fraction of firms only (partial publicity) instead of introducing idiosyncratic noise (partial transparency).

Figure 1. Unconditional Expected Loss and Optimal Degree of Transparency

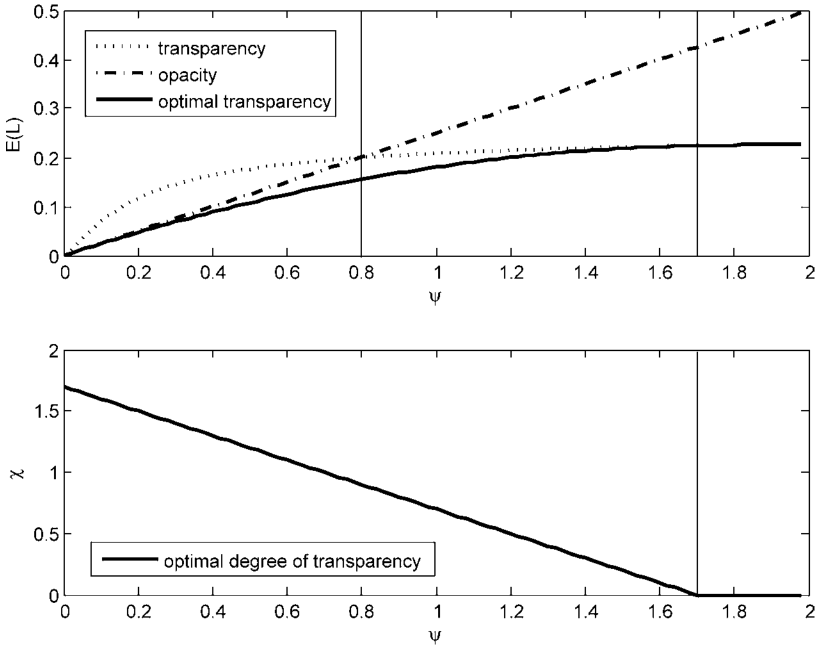

partial transparency when coordination is not very valuable at the social level (λ large), when the degree of strategic complementarities is high (ξ small), and when the relative accuracy of private and central bank information ψ is small.

Figure 1 illustrates the unconditional expected loss under transparency (dotted line), under opacity (dashed line), and under optimal degree of transparency (solid line). The parameter values are $\xi = 0.1$, $\lambda = 1$, and $\sigma_\mu = 0.25$. As (14) shows, full opacity is superior to full transparency when $\psi < \lambda - 2\xi = 0.8$. The optimal degree of transparency is represented in the lower graph shown in figure 1. As (16) states it, reducing the degree of transparency is optimal when $\psi < 2\lambda - 3\xi = 1.7$.

Interestingly, there is a linear negative relation between the relative idiosyncratic private noise ψ and the relative idiosyncratic noise of the central bank's disclosure χ_{exo}. Expression (16) shows that the optimal idiosyncratic relative noise of the central bank's disclosure χ_{exo} declines as the idiosyncratic relative noise of firms' private information ψ increases. The expression $2\lambda - 3\xi$ can be interpreted as the optimal minimal idiosyncratic relative noise, which depends on the preference and structure of the economy.

The distribution of the idiosyncratic noise across ψ and χ is, however, not irrelevant for welfare. The derivation $\frac{\partial \mathbb{E}(L)}{\partial \psi} = 0 \Leftrightarrow \psi = \max[0; -\frac{(\xi+\chi)^2}{2\lambda-\xi+\chi}]$ shows that increasing the relative inaccuracy of firms' private information ψ is detrimental to welfare unless the weight assigned to stabilization λ, the degree of complementarity $1 - \xi$, and the idiosyncratic relative noise of the central bank's disclosure χ are all low $(2\lambda - \xi + \chi < 0)$. This derivation confirms the finding of Hellwig (2005) that increasing the accuracy of firms' private information does not always improve welfare when coordination is significant because more accurate private information would exacerbate dispersion.[12]

Reducing the degree of transparency can be interpreted as the central bank's attempt to control the equilibrium level of coordination.[13] If the central bank could choose the degree of complementarity $1 - \xi$ that drives the coordination motive in the equilibrium pricing rule (12), it would solve $\frac{\partial \mathbb{E}(L)}{\partial \xi} = 0 \Leftrightarrow \xi = \lambda$. The optimal degree of coordination decreases with social aversion to distortion. For instance, when the degree of complementarity is larger than the optimal level for the central bank, i.e., $\xi < \lambda$, the central bank could shape the equilibrium response of firms γ_1 and γ_2 in (12) through the idiosyncratic noise of its disclosure σ_ϕ^2 in the same way that it would by directly influencing ξ. However, the idiosyncratic noise entails a detrimental dispersion effect that mitigates the central bank's incentive to reduce the equilibrium degree of coordination. Therefore, the central bank will reduce the equilibrium degree of coordination less by varying σ_ϕ^2, as it would by varying ξ. From (16), it follows that the central bank will never reduce the degree of transparency when $\xi > \frac{2}{3}\lambda$.

[12]Section 5 discusses the case with microfounded preferences.

[13]See Angeletos and Pavan (2007).

4. Endogenous Central Bank Information

In this section, we drop the assumption that the central bank directly observes the exogenous aggregate shock g that underlies the economy. The central bank has no direct source of information about stochastic aggregate economic conditions, and therefore it must observe the aggregate activity of firms to infer the demand shock. In reality, a central bank learns about aggregate shocks by collecting data from the aggregate economic outcome, not by observing an exogenous fundamental process.[14] The model developed in the current section accounts for the insight that the central bank gains on economic conditions from watching the economy itself. However, the central bank, as a policymaker, also strongly shapes market expectations and thereby drives the course of economic activity. The dual role of the central bank causes a dilemma: the better the central bank succeeds in influencing economic activity, the more the economy reflects the central bank's assessments. Thus, aggregate economic outcomes become less accurate indicators of imbalances.

The next sections describe the information structure and discuss the effect of the central bank's disclosure on the accuracy of its information and of that of firms. Finally, we examine the optimal disclosure strategy and compare it with the case of exogenous information.

4.1 Information Structure

The central bank has no direct access to information on the underlying economic shock. In particular, it cannot observe the aggregate demand shock g. Instead, the central bank bases its estimation of the demand shock on its observations of the overall price level. As Hayek (1945) pointed out, prices play a crucial informational role because they aggregate individual information. By observing the average action of firms, the central bank obtains information about the state of the economy.

[14] In this sense, the realism of the exogenous information structure discussed in section 3 is limited since the signal received by the central bank about economic conditions is independent from the existence or behavior of firms.

We postulate that the central bank receives a signal D on the price level p with some noise η

$$D = p + \eta, \quad \text{with } \eta \sim N(0, \sigma_\eta^2).$$

Using the price-setting rule (12) and the fact that $\gamma_1 + \gamma_2 = 1$, we can express the information of the central bank as

$$D = \gamma_1 g + \gamma_2 D + \eta = g + \frac{\eta}{\gamma_1}.$$

It is important to stress here that the central bank cannot infer the true demand shock g from its observation of the price level D, even if the bank knows which signal it discloses to firms (i.e., D itself), because its observation contains an unknown error η. Therefore, the best central bank estimate of the demand shock, which is conditional on its observation, is the observation itself because the demand shock is improperly distributed: $\mathbb{E}(g|D) = D$.

According to the definition (8), the variance of the central bank expectation error under endogenous information becomes

$$\text{Var}[\mathbb{E}(g|D) - g] \equiv \sigma_\mu^2 = \frac{\sigma_\eta^2}{\gamma_1^2}. \tag{17}$$

As it is immediately visible, the precision of central bank information is a function of the equilibrium response of firms γ_1, which depends—as shown in (12)—upon the disclosure strategy of the central bank.

4.2 Information Value of Prices

Figure 2 illustrates the accuracy of the price level as an indicator of economic conditions. The computation is done with $\sigma_\eta^2 = \sigma_\varepsilon^2 = 0.25$. The information value of prices is evaluated as the variance of the error of demand shock expectations conditional on the price level p. The information value of the price level p is given by

$$\text{Var}[\mathbb{E}(g|p) - g] = \text{Var}\left[\mathbb{E}\left(g|g + \frac{\gamma_2}{\gamma_1}\eta\right) - g\right] = \gamma_2^2 \frac{\sigma_\eta^2}{\gamma_1^2}.$$

The figure shows that the information about the state of the economy that is contained in the price level decreases with the degree of

Figure 2. Informative Value of Prices

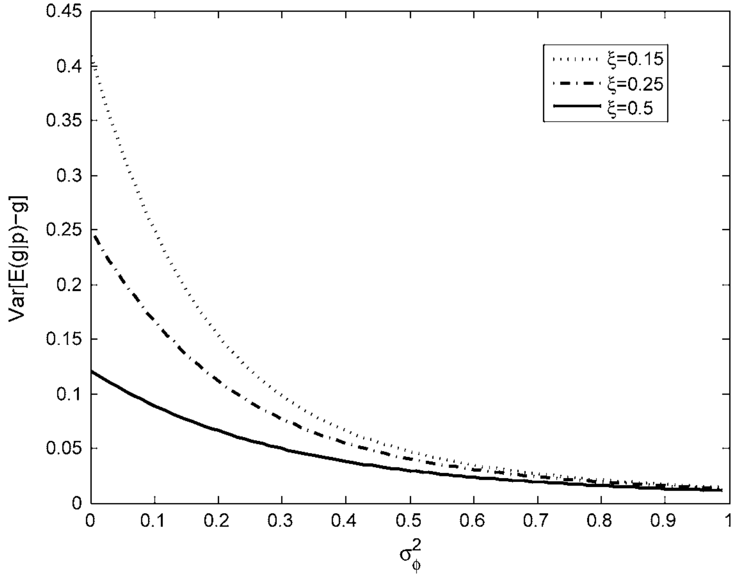

central bank transparency. This relationship clearly highlights the endogenous nature of central bank information. The more effectively the central bank influences the pricing behavior of firms (γ_2 large), the less accurately price levels indicate economic conditions. At the limit of opacity, γ_2 converges to zero (i.e., firms do not react to the disclosure), and the price level becomes a perfect indicator for the demand shock g. The degree of strategic complementarities affects the information value of prices because it drives the overreaction to the central bank's disclosure. When complementarities are high (ξ small), the central bank's disclosure is given a lot of weight in the pricing rule, and this response increases the impact of the noise η on the price level.

4.3 Firms' Information

We now turn to the accuracy of individual firms' information. Although opacity increases the precision of the central bank's observation and thereby of the average disclosure D, it simultaneously increases the idiosyncratic noise (σ_ϕ^2). Opacity therefore has an

Figure 3. Precision of Firms' Information

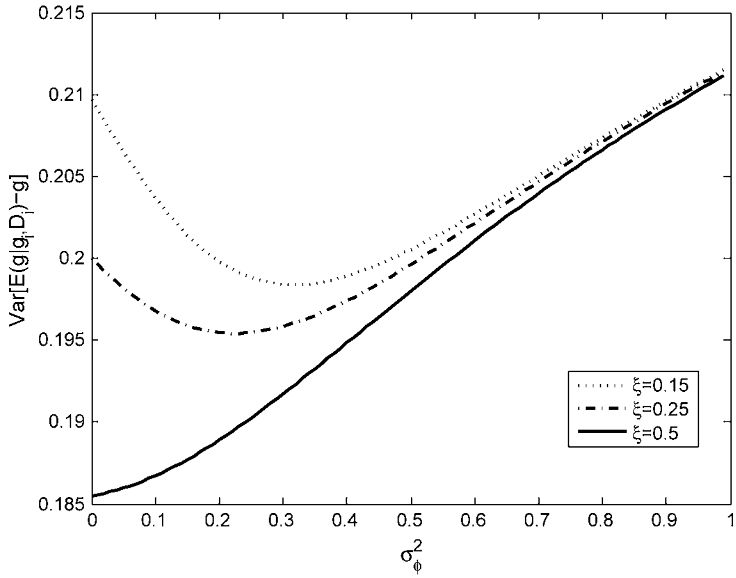

ambiguous impact, overall, on the accuracy of the disclosure that is received by an individual firm i. On one hand, a rise in opacity increases the precision of the central bank's observation and average disclosure. This increase in precision also tends to increase the precision of individual disclosure D_i. On the other hand, a rise in opacity induces a larger idiosyncratic noise ϕ_i that reduces the precision of the individual disclosure. Increasing idiosyncratic noise, which always reduces the precision of firms' information in the case of exogenous information, may, in contrast, increase the precision of firms' information with endogenous information.

The precision of firms' information is interpreted as the variance of the error of expectations for demand shock, conditional on both the private signal g_i and the central bank's disclosure D_i. This is given by

$$\mathrm{Var}[\mathbb{E}(g|g_i, D_i) - g] = \frac{\sigma_\varepsilon^2[\gamma_1^{-2}\sigma_\eta^2 + \sigma_\phi^2]}{\sigma_\varepsilon^2 + \gamma_1^{-2}\sigma_\eta^2 + \sigma_\phi^2}.$$

Figure 3 illustrates this variance for three degrees of strategic complementarities as a function of the degree of transparency. The solid line shows that when complementarities are low, reducing the degree of transparency always decreases firms' information. This decrease occurs because transparency does not significantly distort central bank information when the coordination motive is weak. The increased idiosyncratic noise of opacity is not overcome by the increased precision of the central bank's information.

The dotted and dashed lines show that, when the degree of transparency is high and complementarities are strong, reducing the degree of transparency increases the precision of firms' information. The rise in the precision of central bank information overcomes the rise in idiosyncratic noise as long as transparency is sufficiently high. Below a certain threshold of transparency, lowering transparency further reduces the precision of firms' information. The case of endogenous information highlights a new effect of transparency on the first-order expectation of firms.

4.4 Welfare

Given the definition of the variance of the central bank expectation error under endogenous information (17) and the equilibrium behavior of firms (12), the unconditional expected social loss (7) can be written as

$$\mathbb{E}(L) = \gamma_1^2 \sigma_\varepsilon^2 + \gamma_2^2 \sigma_\phi^2 + \lambda \gamma_2^2 \frac{\sigma_\eta^2}{\gamma_1^2}. \tag{18}$$

4.5 Optimal Degree of Transparency

This section derives the optimal degree of transparency under endogenous information and compares it with the case of exogenous information discussed in section 3.4.

To determine the optimal degree of transparency σ_ϕ^2, we derive the loss (18) with respect to σ_ϕ^2 and set it equal to zero:

$$\begin{aligned}\frac{\partial \mathbb{E}(L)}{\partial \sigma_\phi^2} &= 2\gamma_1 \sigma_\varepsilon^2 \frac{\partial \gamma_1}{\partial \sigma_\phi^2} + \gamma_2^2 + 2\gamma_2 \sigma_\phi^2 \frac{\partial \gamma_2}{\partial \sigma_\phi^2} + 2\lambda\gamma_2 \frac{\sigma_\eta^2}{\gamma_1^2} \frac{\partial \gamma_2}{\partial \sigma_\phi^2} \underbrace{-2\lambda\gamma_2^2 \frac{\sigma_\eta^2}{\gamma_1^3} \frac{\partial \gamma_1}{\partial \sigma_\phi^2}}_{\text{endog. spec.}} \\ &= 0. \end{aligned} \tag{19}$$

The last term of this equation captures the endogenous specific impact that the central bank's transparency has on the accuracy of firms' information. Solving this optimization problem analytically is not straightforward because the equilibrium pricing rule (12) is characterized by non-linear equations under endogenous information. However, it is easy to recognize two properties of the optimal degree of transparency for endogenous information, which is interpreted as the optimal relative variance of idiosyncratic and fundamental noise associated with the central bank's disclosure $\chi_{\text{endo}} = \frac{\sigma_\phi^2 \gamma_1^2}{\sigma_\eta^2}$. First, the optimal degree of transparency under exogenous information χ_{exo} converges to that under endogenous information χ_{endo} when the relative inaccuracy of private and central bank information ψ approaches zero.[15] Second, the optimal degree of transparency under endogenous information χ_{endo} is unrelated to the relative inaccuracy of private and central bank information ψ. These observations allow us to analytically express[16] the optimal degree of transparency under endogenous information as

$$\chi_{\text{endo}} = \max[0; 2\lambda - 3\xi]. \tag{20}$$

Reducing transparency increases the accuracy of firms' average information under endogenous information, and therefore it is not surprising that the optimal degree of transparency under endogenous information is lower than it is under exogenous information: $\chi_{\text{endo}} \geq \chi_{\text{exo}}$ (compare with (16)). Contrary to exogenous information, the optimal degree of transparency under endogenous information is independent of the relative inaccuracy of private and central bank information ψ. The central bank does not find it optimal to increase the degree of transparency in response to less accurate firms' information because the accuracy of its information depends on firms' reaction to its disclosure. An increase in the inaccuracy of firms' private information makes the central bank's disclosure relatively more accurate, which would incite the central bank to become more transparent. However, a higher degree of transparency would decrease the accuracy of central bank disclosure. It turns out that

[15] More precisely, this occurs when $\psi_{\text{exo}} \equiv \frac{\sigma_\varepsilon^2}{\sigma_\mu^2} = \psi_{\text{endo}} \equiv \frac{\sigma_\varepsilon^2 \gamma_1^2}{\sigma_\eta^2}$ approaches zero.

[16] Numerical procedures unambiguously confirm this analytical expression.

Figure 4. Optimal Degree of Transparency for Exogenous and Endogenous Information

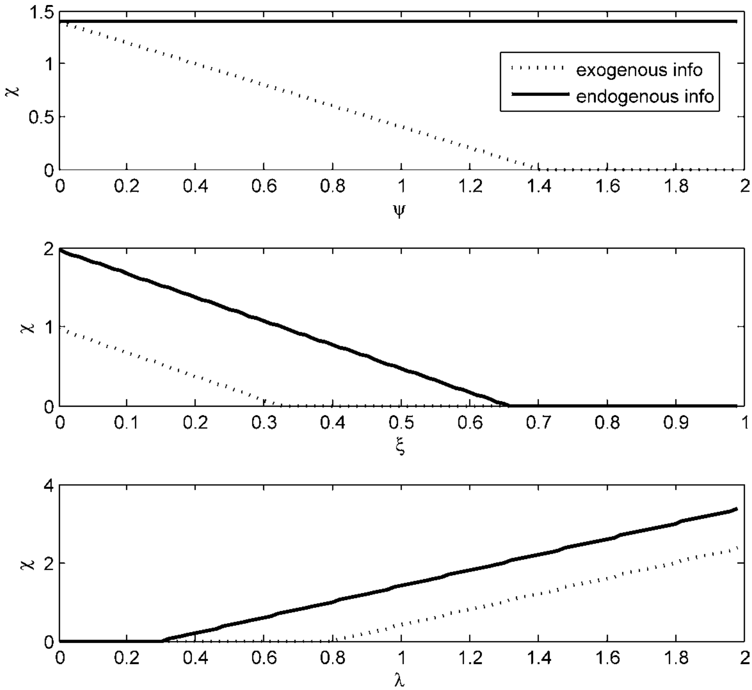

both effects perfectly offset each other. With endogenous information, the optimal minimal idiosyncratic relative noise $2\lambda - 3\xi$ exclusively impairs the central bank's disclosure and cannot be traded off against the idiosyncratic noise of firms' private information, as it is with exogenous information.

Figure 4 illustrates the optimal degree of transparency with exogenous (dotted line) and endogenous (solid line) central bank information. The parameter values set by default are $\psi = 1$, $\xi = 0.25$, and $\lambda = 1$. The first graph illustrates the impact of ψ, the relative accuracy of private and central bank information, on the optimal degree of transparency. Under exogenous information, an increase in the idiosyncratic relative noise of firms' private information ψ implies a reduction of the idiosyncratic relative noise of central bank disclosure χ (an increase in the degree of transparency). By contrast, an increase in ψ does not alter the optimal degree of transparency under endogenous information. The second and third graphs illustrate the optimal degree of transparency as a function of the degree of complementarities $1 - \xi$ and the weight assigned to the output distortion λ, which determine the optimal minimal idiosyncratic relative noise.

5. Microfounded Welfare Preferences

This section examines the welfare effect of transparency when the weight assigned to output-gap deviation λ is consistent with the other parameters of the model. Adam (2007) derived the microfounded welfare of the representative household and showed that $\lambda = \frac{\xi}{\theta}$, where $\theta > 1$ is the price elasticity of demand in the Dixit-Stiglitz aggregator. The loss function (7) then becomes

$$L = \int_i (p_i - p)^2 di + \frac{\xi}{\theta}(g - p)^2. \tag{21}$$

The relative weight assigned to coordination is high because price dispersion reduces the utility of the representative household. This increases the welfare-improving effect of the central bank's disclosure because it helps to coordinate firms' pricing decisions. We may therefore expect the optimal degree of transparency to be higher than it is when coordination and stabilization are equally weighted, as in M-S ($\lambda = 1$).

Because the optimal minimal idiosyncratic relative noise $2\lambda - 3\xi = 2\frac{\xi}{\theta} - 3\xi < 0$ is negative with the microfounded weight λ, we conclude that full transparency is the best disclosure strategy under both exogenous and endogenous information.

It should be noted that the model with exogenous information confirms the study of Hellwig (2005), which analyzed the welfare effect of public disclosure in a fully microfounded model. Under exogenous information, the unconditional expected loss steadily increases with the inaccuracy of central bank information $\frac{\partial \mathbb{E}(L)}{\partial \sigma_\mu^2} > 0$. However, reducing the relative inaccuracy of private and central bank information ψ does not always improve welfare: $\frac{\partial \mathbb{E}(L)}{\partial \psi} > 0 \Leftrightarrow \psi > \frac{\xi\theta}{\theta - 2}$. Coordination is highly valuable for microfounded preferences; therefore, increasing the accuracy of firms' private information may reduce welfare because price dispersion will be exacerbated.

6. Concluding Remarks

This paper developed a model for endogenous central bank information that creates a trade-off between shaping market expectations and learning from them. As transparency increases the effect of central bank disclosure on firms' behavior, it reduces the accuracy of prices as an indicator for underlying shocks and thereby the accuracy of central bank information.

This study found that the endogeneity of central bank information entails an additional detrimental effect of transparency to that highlighted by Morris and Shin (2002). Their model of exogenous information underlined the potentially detrimental effect of transparency when higher-order expectations are given a large weight because of coordination. However, under endogenous information, transparency also decreases the accuracy of central bank information, which affects both the first- and higher-order expectations of firms.

We can draw two conclusions related to welfare. First, accounting for the endogeneity of information highlights the detrimental effect of transparency. The optimal degree of transparency is always lower under endogenous information than under exogenous information. Second, the optimal degree of transparency is unrelated to the relative accuracy of firms' private information under endogenous information. Thus, whether central bank information is more accurate than the information that is privately available in the economy is irrelevant to address central bank transparency questions. Instead, the optimal degree of transparency depends on the preference for coordination versus stabilization and the degree of complementarity in the economy.

References

Adam, K. 2007. "Optimal Monetary Policy with Imperfect Common Knowledge." *Journal of Monetary Economics* 54 (2): 267–301.

Amato, J., S. Morris, and H. S. Shin. 2002. "Communication and Monetary Policy." *Oxford Review of Economic Policy* 18 (4): 495–503.

Amato, J., and H. S. Shin. 2006. "Imperfect Common Knowledge and the Information Value of Prices." *Economic Theory* 27 (1): 213–41.

Angeletos, G.-M., and A. Pavan. 2007. "Efficient Use of Information and Social Value of Information." *Econometrica* 75 (4): 1103–42.

Baeriswyl, R., and C. Cornand. 2010. "The Signaling Role of Policy Actions." *Journal of Monetary Economics* 57 (6): 682–95.

Bernanke, B. S., and M. Woodford. 1997. "Inflation Forecasts and Monetary Policy." *Journal of Money, Credit, and Banking* 29 (4): 653–84.

Cornand, C., and F. Heinemann. 2008. "Optimal Degree of Public Information Dissemination." *Economic Journal* 118 (528): 718–42.

Cukierman, A., and A. H. Meltzer. 1986. "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information." *Econometrica* 54 (5): 1099–1128.

Ehrmann, M., and M. Fratzscher. 2005. "Transparency, Disclosure and the Federal Reserve." ECB Working Paper No. 457.

Faust, J., and L. E. O. Svensson. 2002. "The Equilibrium Degree of Transparency and Control in Monetary Policy." *Journal of Money, Credit, and Banking* 34 (2): 520–39.

Geithner, T. F. 2006. "Uncertainty and Transparency in the Conduct of Monetary Policy." *Review* (Bank for International Settlements) 54.

Geraats, P. M. 2002. "Central Bank Transparency." *Economic Journal* 112 (483): F532–65.

———. 2009. "Trends in Monetary Policy Transparency." *International Finance* 12 (2): 235–68.

Hayek, F. 1945. "The Use of Knowledge in Society." *American Economic Review* 35: 519–30.

Heinemann, F., and G. Illing. 2002. "Speculative Attacks: Unique Equilibrium and Transparency." *Journal of International Economics* 58 (2): 429–50.

Hellwig, C. 2005. "Heterogeneous Information and the Welfare Effects of Public Information Disclosures." Mimeo. Available at www.econ.ucla.edu/people/papers/Hellwig/Hellwig283.pdf.

Hellwig, C., and L. Veldkamp. 2009. "Knowing What Others Know: Coordination Motives in Information Acquisition." *Review of Economic Studies* 76 (1): 223–51.

Morris, S., and H. S. Shin. 2002. "Social Value of Public Information." *American Economic Review* 92 (5): 1521–34.

———. 2005. "Central Bank Transparency and the Signal Value of Prices." *Brookings Papers on Economic Activity* 36 (2): 1–66.

Romer, C. D., and D. H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review* 90 (3): 429–57.

Svensson, L. E. O. 2006. "Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro-Transparency, Not Con." *American Economic Review* 96 (1): 448–51.

Walsh, C. E. 2007. "Optimal Economic Transparency." *International Journal of Central Banking* 3 (1): 5–36.

Woodford, M. 2005. "Central Bank Communication and Policy Effectiveness." NBER Working Paper No. 11898.

Inflation Risk Premia and Survey Evidence on Macroeconomic Uncertainty*

Paul Söderlind
University of St. Gallen and CEPR

The difference between nominal and real interest rates (break-even inflation) is often used to gauge the market's inflation expectations—and has become an important tool in monetary policy analysis. However, break-even inflation can move in response to shifts in inflation risk premia and liquidity premia as well as to changes in expected inflation. This paper sheds light on this issue by analyzing the evolution of U.S. break-even inflation from 1997 to mid-2008. Regression results show that survey data on inflation uncertainty and proxies for liquidity premia are important factors.

JEL Codes: E27, E47.

1. Introduction

Break-even inflation (the difference between nominal and real interest rates) is often used as an indicator of inflation expectations. It has several advantages over surveys and other methods: It is available on a daily basis, it focuses on the beliefs of financial markets, and it is based on decisions that matter ("put your money where your mouth is"). On the other hand, survey data is often considered to give a cleaner measure of inflation expectations. It is therefore important to study how the break-even inflation rates differ from survey data on inflation expectations—and to find proxies that can account for the difference. In addition, understanding the pricing of Treasury bonds is important for both analysts and investors.

*I thank Michael Fischer, Angelo Ranaldo, the editor Andrew Levin, and an anonymous referee for comments; Juan Angel García and Andrés Manzanares for fruitful discussions at an early stage of this work; and Kenneth F. Wallis for providing extended data from one of his articles. Author address: SBF, University of St. Gallen, Rosenbergstrasse 52, CH-9000 St. Gallen, Switzerland. E-mail: Paul.Soderlind@unisg.ch.

The idea of this paper is very simple. Data on the U.S. nominal and real interest rates is combined with survey data on inflation expectations to construct a time series of "break-even deviations"—that is, the deviation of break-even inflation from inflation expectations—for the sample 1997:M1–2008:M6. This deviation is likely to be driven mainly by inflation risk premia (on the nominal bonds) and liquidity premia (on the real bonds). It is therefore regressed on measures of inflation uncertainty from survey data on probability distributions and proxies for real liquidity premia. The results indicate that the regressors are significant and explain a considerable fraction of the movements of the break-even deviation. This leads to significant adjustments of the usual break-even inflation.

Several studies have looked at similar issues, but very few have combined real interest rates, survey data on inflation expectations and inflation uncertainty, and proxies for real liquidity premia to estimate a Fisher equation. (However, recent work by García and Werner 2010 uses a similar approach on European Central Bank [ECB] data.)

Reschreiter (2004), Chen, Liu, and Cheng (2005), and Grishchenko and Huang (2008) use nominal and real interest rates to estimate no-arbitrage yield-curve models with latent factors. D'Amico, Kim, and Wei (2008) and Joyce, Lildholdt, and Sorensen (2009) go one step further by including also inflation expectations. The current paper does not estimate a fully fledged yield-curve model (mostly because inflation expectations are only available for a few horizons), but it uses observable factors to capture the inflation risk and real liquidity premia. This might facilitate the interpretation of the results.

Lahiri, Teigland, and Zaporowski (1988), García and Manzanares (2007), and Wright (2008) also use survey data to estimate inflation uncertainty. The first paper does not find any effect of uncertainty on nominal interest rates (on the sample 1969–86), but the second paper finds that it is important for understanding the inflation scares during the Volcker period in the early 1980s. The current paper differs from these earlier contributions in many details—in particular, on how to estimate uncertainty from survey data, but more importantly by focusing on the period when real (inflation-indexed) interest rates are available (since 1997). This allows for controlling for an important driving force of nominal rates.

Carlstrom and Fuerst (2004) and Shen (2006) compare the break-even inflation rate with survey measures of inflation expectations. They argue that the break-even inflation rates are too low—due to a considerable liquidity premium in the observed real rates. The current paper incorporates those results—and goes on to study if other factors (inflation uncertainty/disagreement) add explanatory value by capturing the inflation risk premium. Kajuth and Watzka (2008) is similar, but does not use a survey-based direct measure of inflation uncertainty (instead, they study inflation disagreement).

The current paper draws on all these strands of research in order to estimate a modern Fisher equation where we use information on real interest rates and inflation expectations but also control for the liquidity premium on real bonds and the inflation risk premium on nominal bonds. The liquidity premium is modeled in terms of observed liquidity premia for off-the-run Treasury bonds (and some other proxies) and the inflation risk premium by unique information about inflation risk reported by the participants of the Survey of Professional Forecasters (and some other proxies). The results indicate that these variables are important—and lead to significant adjustments of break-even inflation.

The outline of the paper is as follows. Section 2 discusses the Fisher equation and the data, section 3 presents the empirical results, section 4 shows break-even inflation rates and inflation uncertainty for other economies (the United Kingdom and the euro area), and section 5 concludes.

2. Data and Relation to the Fisher Equation

This section summarizes the Fisher equation, and describes the data and how to construct proxies for inflation uncertainty and liquidity premia.

2.1 The Fisher Equation

A modern Fisher equation says that the nominal interest rate (i) equals the sum of expected inflation ($\mathrm{E}\pi$), the real interest rate ($\tilde{r}$), and an inflation risk premium (IRP):

$$i = \mathrm{E}\pi + \tilde{r} + IRP. \tag{1}$$

(For notational simplicity, subscripts for time and maturity are suppressed.)

However, the real (inflation-indexed) bonds traded in the United States have occasionally been somewhat illiquid, and it is often argued that they carry a liquidity premium (see Shen 2006 and also Roush 2008 for a detailed discussion of the TIPS market). Therefore, let r indicate the real interest rate observed in data and let LPr indicate the liquidity premium on real bonds. The observed real interest rate needs to be adjusted down ($\tilde{r} = r - LPr$) to correct for this liquidity premium. The Fisher equation for observable data can then be written

$$i - r - \mathrm{E}\pi = IRP - LPr. \qquad (2)$$

The observed "break-even deviation" ($i - r - \mathrm{E}\pi$) depends positively on the inflation risk premium and negatively on the liquidity premium on real bonds. The break-even deviation is likely to be stationary, even if nominal interest rates and inflation expectations are not.

Central banks often calculate break-even inflation as the difference between nominal and real yields ($i - r$)—with the purpose of approximating the market's inflation expectations. If the inflation risk premium and the liquidity premium in the Fisher equation (2) are constant, then the break-even inflation rate differs from the inflation expectations by a constant: changes in the break-even inflation rate are then precise measures of changes in inflation expectations. Otherwise, an increased inflation risk premium increases the break-even inflation rate, while an increased real liquidity premium decreases it—and these movements may hurt the informational value of the break-even inflation rate.

The objective of this paper is to find proxies for the inflation risk premium and the real liquidity premium in order to understand the pricing of bonds and to improve upon the practice of calculating break-even inflation.

2.2 Data

The regressions estimated below are based on equation (2), and this section discusses how the variables are constructed.

Figure 1. Interest Rates (Nominal and Real) and Inflation Expectations

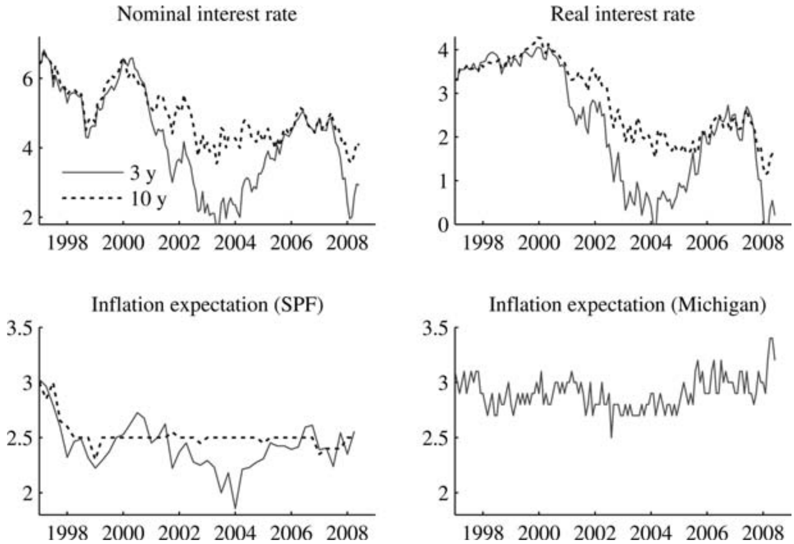

Note: This figure shows estimated nominal and real zero-coupon rates from McCulloch's home page, as well as inflation expectations from the SPF and the Michigan survey.

The three- and ten-year nominal and real (zero-coupon) interest rates from McCulloch (2008) are shown in the upper panels of figure 1. I choose to focus mostly on the horizons of three and ten years: the real rates for shorter maturities are less reliable, for instance, because there have occasionally been few outstanding index-linked bonds with short maturities. Other maturities are used in the robustness analysis and briefly discussed in the text. The sample starts in 1997:M1 since this is when the Treasury Inflation-Protected Securities (TIPS) were first issued, and the data is available on a monthly frequency.[1]

The main measures of inflation expectations and inflation risk uncertainty are from the Survey of Professional Forecasters (SPF)

[1] The data from Gürkaynack, Sack, and Wright (2008) are very similar. For instance, the end-of-month break-even inflation rates (for the ten-year horizon) have a correlation of 0.97 with those from McCulloch.

Figure 2. Inflation Uncertainty and Off-the-Run Liquidity Premium

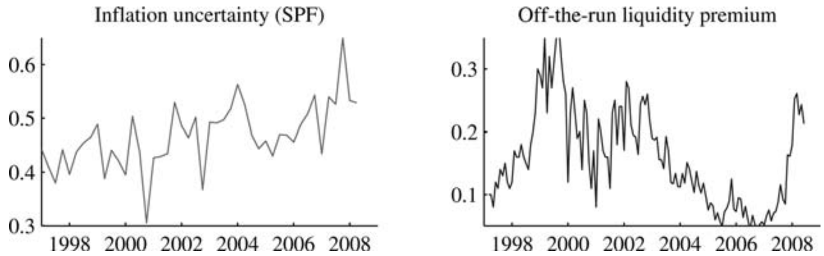

Note: This figure shows the inflation uncertainty from the SPF and the off-the-run liquidity premium from the Federal Reserve Bank of Cleveland.

at the Federal Reserve Bank of Philadelphia (2010). The SPF is a quarterly survey of forecasters' views on key economic variables. The respondents, who supply anonymous answers, are professional forecasters from the business and financial community. The survey asks for, among other things, the next-year and the ten-year CPI inflation forecasts, but also for probability distributions (histograms) of GDP deflator inflation. The deadline for the survey is approximately in the middle of the quarter.[2]

The lower-left panel of figure 1 shows three- and ten-year inflation expectations constructed from the median (across forecasters) CPI inflation forecasts. The three-year expectation is approximated by linear interpolation of the forecasts for the next calendar year and for the next ten years. There are substantial movements in the three-year expectations, while the ten-year inflation expectations are virtually flat after 1999.

The left panel of figure 2 shows inflation uncertainty estimated from the probability distributions of inflation in the SPF. This variable measures the average uncertainty about future inflation—as reported by the participants of the SPF. It has the great advantage of being a direct measure of inflation uncertainty, rather than some proxy (like inflation disagreement/forecast dispersion or implied volatility from bond options).

[2]While deflator inflation and CPI inflation differ in many respects, they tend to be strongly correlated in the medium run. For instance, the correlation of four-quarter inflation is 0.94 on data from 1970 through 2008.

The estimation of inflation uncertainty involves fitting distributions to the histograms in the SPF (there is one histogram per forecaster in each quarter). If only one bin is used (the respondent puts 100 percent of the probability on one of the prespecified bins), then a triangular distribution within that bin is assumed. If two or more bins are used, then a normal distribution is estimated by minimizing the sum of the squared deviations of the theoretical from the observed probabilities. (See Giordani and Söderlind 2003 for an early application and García and Manzanares 2007 for a critique of the least-squares criterion.) Clearly, the analysis of long-term interest rates asks for information about long-run uncertainty. Unfortunately, the SPF contains only information on short-run uncertainty. I am therefore left with using the short-run information—hoping that long-run uncertainty can be approximated by (linear functions of) short-term uncertainty.[3]

These estimations generate a panel of subjective standard deviations of future inflation. However, there is a lot of cross-sectional (across forecasters) dispersion in the estimation results—some of which appears to be caused by typos and other data errors. To get robust estimates of the average individual uncertainty, I use the cross-sectional trimmed mean (20 percent trimming from both bottom and top) of the individual standard deviations. (Using the median gives very similar results and using the mean produces fairly similar, but more erratic, results.)

An additional complication is that the probability distributions are for year-on-year inflation (the value in a calendar year divided by the value in the previous calendar year, minus one). This means that the forecasting horizon varies across the sample: the current-year forecast made in Q1 has a four-quarter horizon, the forecast made in Q2 has a three-quarter horizon, and so forth. The estimated uncertainty therefore has clear seasonality—as the effective forecasting horizon decreases over the calendar year. My measure of inflation uncertainty is therefore deseasonalized using X12 (assuming

[3]The SPF also contains histograms about inflation during the next calendar year, but that seems to convey essentially the same information—except that it appears more noisy. The uncertainty about the next calendar year is of course higher, but the correlation of year 1 and 2 is strong (around 0.75)—and the regression results are similar (not tabulated).

multiplicative seasonality). The resulting series (in figure 2) still appears a bit noisy, so it should come as no surprise that the regression analysis (see below) favors an average of the current and lagged uncertainty as a regressor.

The right panel of figure 2 shows the off-the-run liquidity premium: the yield difference between less and more liquid ("off-the-run" and "on-the-run") ten-year nominal Treasury bonds. The Federal Reserve Bank of Cleveland (2010) calculates an adjusted break-even inflation rate by adding a real liquidity premium (LPr), that is, as $i - (r - LPr)$.[4] To construct the real liquidity premium (which is not directly observable), the Federal Reserve Bank of Cleveland (2010) argues that it is likely to be strongly related to the off-the-run premium according to a linear-quadratic function.[5] This function has a negative constant which is meant to capture a constant inflation risk premium of the nominal bond. For the sort of off-the-run premia that we see in the sample, the function is actually almost linear with a slope of approximately 3.5, except for some rare cases (the peak in 1999, when the quadratic term dampens the effect). The publication of the LPr was discontinued in autumn 2008, as it was deemed unreliable during the financial crisis. Clearly, liquidity premia behaved distinctly different during the crisis, so I choose to end my sample at the same time as the Federal Reserve Bank of Cleveland (in 2008:Q2).[6]

In the main specification, I use the data described above. The additional analysis also considers several other proxies for expected inflation, inflation uncertainty, and the real liquidity premium. This provides a robustness analysis and explores the possibility of using the monthly (instead of quarterly) data presented below.

The lower-right panel of figure 1 shows the median five-year inflation expectations according to the Survey of Consumers at

[4] See also Graveline and McBrady (2006) for a discussion of the off-on yield spread.

[5] The functional form is $LPr = -0.948 + 12.71LPo - 20.9(LPo)^2$, where LPo is the off-the-run premium.

[6] On October 31, 2008, the Federal Reserve Bank of Cleveland (2010) stated: "We have discontinued the liquidity-adjusted TIPS expected inflation estimates for the time being. The adjustment was designed for more normal liquidity premiums. We believe that the extreme rush to liquidity is affecting the accuracy of the estimates."

Figure 3. Alternative Proxies for Inflation Uncertainty and Liquidity Premium

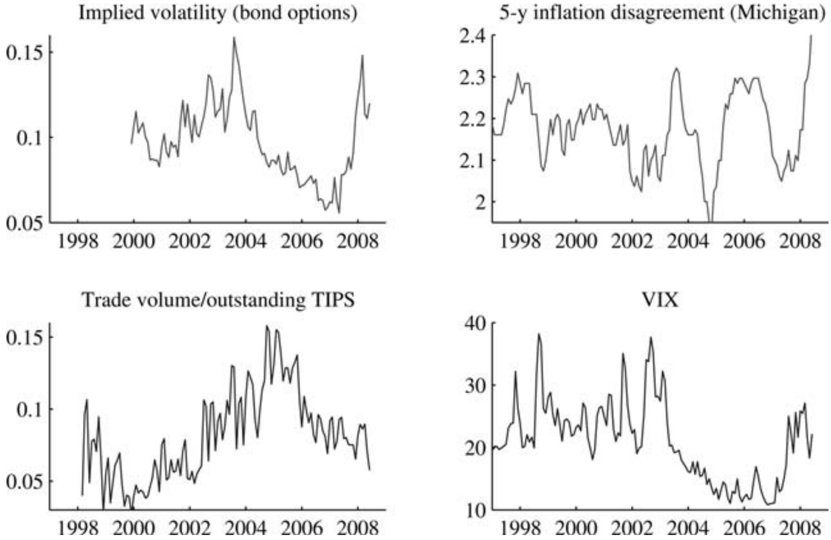

Note: This figure shows the implied volatilities from bond options, the disagreement of five-year inflation forecasts in the Michigan survey, the TIPS trade as a fraction of the outstanding volume, and the VIX.

Reuters/University of Michigan (2010). The Michigan survey is a monthly survey of approximately 500 randomly chosen consumers. The preliminary data is released in the middle of the month. Compared with the SPF, the inflation expectations appear to be somewhat higher (hovering around 3 percent instead of 2.5 percent). However, the general movements appear to be in line with the SPF—with a variability somewhere between that of the three- and ten-year expectations (except, possibly, toward the end of the sample, where the Michigan survey increases substantially).

The upper panels of figure 3 show two alternative proxies for inflation uncertainty. The upper-left subfigure shows (monthly averages of) implied volatilities from CME's options on bond futures.[7]

[7]The futures are on a thirty-year Treasury bond and are available since December 1999. Data for options on bonds with shorter maturities are available only for a shorter sample.

Similar to the SPF inflation uncertainty in figure 2, there are local peaks in 2002, 2004, and toward the end of the sample. The upper-right subfigure shows the disagreement (cross-sectional interquartile range) of the five-year inflation forecasts in the monthly Michigan survey—the idea being that disagreement may be correlated with uncertainty about future inflation (see, for instance, Giordani and Söderlind 2003 for a critical assessment).[8] The disagreement shares some movements with the other measures of inflation risk, but shows some idiosyncratic patterns—for instance, the local peak in 2006. (Similar proxies for inflation uncertainty have previously been used by D'Amico, Kim, and Wei 2008, Gürkaynak, Sack, and Wright 2008, and Ciccarelli and García 2009.)

The lower panels of figure 3 show two alternative (monthly) proxies for the liquidity premium on real bonds: the TIPS trade volume from the Federal Reserve Bank of New York (2010) in relation to the volume of outstanding TIPS from the U.S. Treasury (2010) and the Chicago Board of Exchanges's volatility index VIX (derived from options on the S&P 500). Both variables have some similarities with the off-the-run premium (with a reversed sign for the VIX).[9]

3. Empirical Results

3.1 Long-Term Break-Even Inflation

Figure 4 (left panel) shows the "break-even deviation" ($i - r - \mathrm{E}\pi$) for the ten- and three-year horizons—which are the dependent variables in the regressions below (except in some of the subsequent robustness analysis). These deviations peak in 1997, 2000, and 2004–08, with troughs in 1999 and 2002–03.

Table 1 reports results from a number of regressions of the ten-year break-even deviation based on model (2). The sample is quarterly (since the SPF is) and covers the period 1997:Q1–2008:Q2.

The first column shows that the off-the-run premium has a coefficient of −3.07, which is close to the slope in the Federal Reserve Bank of Cleveland (2010) discussed above, and that the inflation

[8] It can be shown that the disagreement about one-year inflation gives similar results.

[9] See Fleming and Krishnan (2009) for a discussion of the TIPS market.

Figure 4. "Break-Even Deviation" and Break-Even Inflation

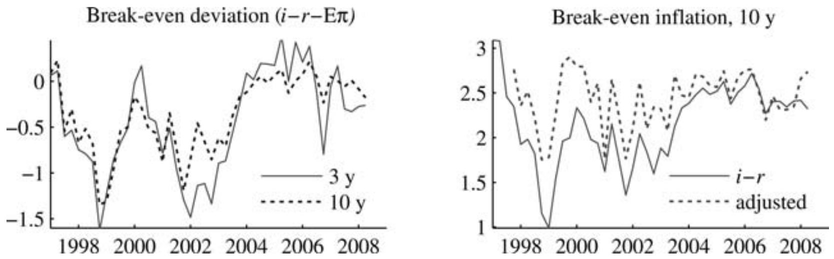

Note: This figure shows the break-even deviation (nominal interest rate minus the real interest rate minus the inflation expectations according to the SPF) and the break-even inflation (nominal minus real interest rate, with/without adjustment).

Table 1. Regression Results, Break-Even Deviation for the Ten-Year Maturity, Quarterly Data

	(1)	(2)	(3)	(4)	(5)
Inflation Uncertainty	2.77	2.38	1.94	2.36	3.19
	(3.81)	(3.02)	(2.56)	(2.81)	(4.24)
Implied Volatility				−3.81	
				(−1.22)	
Inflation Disagreement					0.69
					(1.48)
Off-the-Run Premium	−3.07	−2.88	−0.79	−2.18	−2.95
	(−3.68)	(−3.43)	(−1.27)	(−2.03)	(−3.70)
TIPS Trade		1.59			
		(1.04)			
VIX			−0.04		
			(−5.81)		
Constant	−0.38	−0.37	−0.38	−0.34	−0.38
	(−7.33)	(−7.00)	(−9.82)	(−6.59)	(−7.60)
R^2	0.45	0.46	0.67	0.42	0.47
Observations	43.00	41.00	43.00	34.00	43.00

Notes: The table shows regression coefficients and t-statistics (in parentheses) for quarterly data 1997:Q1–2008:Q2. The dependent variable is the nominal interest rate minus the real interest rate and minus expected inflation (ten-year horizon). All regressors (except the constant) are demeaned. The t-statistics are based on a Newey-West estimator with one lag.

uncertainty has a coefficient of 2.77. According to Newey-West standard errors, both coefficients are significant at any traditional significance level. Based on the signs of the coefficients, these variables can be thought of as capturing the real liquidity premium and the inflation risk premium, respectively. In this and all subsequent regressions, "inflation uncertainty" is measured as the moving average over the current and lagged quarter. The reason appears to be that quarterly data is fairly noisy, but the broader movements of subjective beliefs about future inflation are strongly related to the break-even deviation.

According to these estimates, the real liquidity premium and the inflation risk premia contribute significantly to the movements of the break-even deviation seen in figure 4. For instance, the off-the-run premium and the inflation uncertainty both have a range of 0.2 (see figure 2). Multiplying by 3 (the coefficients are −3.07 and 2.77, respectively) suggests that each component moves the fitted values by 0.6, which is around a third of the observed movements of the break-even deviation.

To be more precise, the trough of the break-even deviation in 1999 (see figure 4) is mostly due to a high real liquidity premium: the off-the-run liquidity premium is high then (perhaps related to the Y2K scare). The real liquidity premium falls back over the next year, creating a local peak in the break-even deviation in 2000. The moderately low break-even deviation in 2002–03 (corresponding roughly to the period of deflation scare) is similar to the trough in 1999 (that is, mostly driven by a high liquidity premium), except that the increase in the inflation risk premium prevents the break-even deviation from falling very much. The very high break-even deviation in 2004–07 is initially driven by a low liquidity premium—and later increased further by a higher inflation risk premium.

The implications for break-even inflation are shown in the right panel of figure 4. By reshuffling (2) we have

$$i - r + LPr - IRP = \mathrm{E}\pi. \tag{3}$$

The figure shows both the traditional break-even inflation rate $(i-r)$ and an adjusted one $(i - r + LPr - IRP)$. The adjustment amounts to making the troughs in 1999 and 2002–03 much less pronounced but the peak in 2000 higher.

The policy relevance of these findings is hard to judge. Clearly, policy analysis and decisions are based on a very large information set, and break-even inflation is only one piece. However, if we were to entertain the idea that break-even inflation is a key number, then the findings in figure 4 have important implications. In particular, they suggest that monetary policy should have been tighter during much of the early part of the sample.

The other four columns of table 1 provide variations on the same specification by adding other proxies. It is worth noticing that inflation uncertainty stays significantly positive in all variations and that the off-the-run premium is significantly negative in all but one case.

Columns 2 and 3 add alternative proxies for the real liquidity premium. In column 2, the trading volume of TIPS (as a fraction of the outstanding volume) has a positive point estimate (indicating a negative relation with the liquidity premium of the real interest rate), but it is far from being statistically significant. In column 3, the VIX has a negative (and strongly significant) coefficient—suggesting that the liquidity premium of the real interest rate increases when equity market volatility does. In fact, it seems as if the VIX dominates the off-the-run premium in accounting for the real liquidity premium (without destroying the importance of the inflation risk variable). Still, these variables seem to be close substitutes—and the off-the-run premium is by now a reasonably well-established proxy (see the Federal Reserve Bank of Cleveland 2010) and has the appeal of focusing on the liquidity on the Treasury bond market, rather than general financial turbulence.

Columns 4 and 5 of table 1 add alternative proxies for inflation uncertainty. In column 4, the implied volatility from bond options shows up with the wrong sign and is far from significant. In column 5, the disagreement (measured as the cross-sectional interquartile range) among the respondents of the survey has the expected positive sign, but it is also not significant. It can also be shown (not tabulated) that uncertainty of future output growth (according to the distributions reported to the SPF) is not significant.

Overall, table 1 suggests a simple specification where the inflation risk premium is approximated by the inflation uncertainty from the SPF and the real liquidity premium is approximated by either the off-the-run premium or the VIX. All other proxies are insignificant.

Table 2. Regression Results, Break-Even Deviation for the Three-Year Maturity, Quarterly Data

	(1)	(2)	(3)	(4)	(5)
Inflation Uncertainty	2.38	1.65	1.23	1.61	3.18
	(1.93)	(1.28)	(1.16)	(1.12)	(2.62)
Implied Volatility				−2.67	
				(−0.54)	
Inflation Disagreement					1.32
					(1.36)
Off-the-Run Premium	−3.85	−3.46	−0.71	−4.10	−3.62
	(−3.27)	(−3.02)	(−0.94)	(−1.82)	(−3.24)
TIPS Trade		3.12			
		(1.00)			
VIX			−0.05		
			(−7.28)		
Constant	−0.45	−0.44	−0.45	−0.42	−0.45
	(−5.17)	(−4.94)	(−6.70)	(−3.77)	(−5.36)
R^2	0.32	0.34	0.54	0.30	0.36
Observations	43.00	41.00	43.00	34.00	43.00

Notes: The table shows regression coefficients and *t*-statistics (in parentheses) for quarterly data 1997:Q1–2008:Q2. The dependent variable is the nominal interest rate minus the real interest rate and minus expected inflation (three-year horizon). All regressors (except the constant) are demeaned. The *t*-statistics are based on a Newey-West estimator with one lag.

3.2 Comparison across Maturities

Table 2 is similar to table 1, except that it studies the three-year horizon. The point estimates are very similar to those for the ten-year horizon, but the significance of inflation uncertainty is weaker (the liquidity premium and the VIX are still strongly significant).

Results for other maturities (not tabulated) indicate that the five-year horizon looks very similar to the ten-year horizon (this is also true for the five-year "forward rate" starting five years ahead). In contrast, the one-year horizon (with all the caveats about the quality of the TIPS data for short maturities) is fairly similar to the three-year horizon. Overall, it seems as if liquidity premia and the VIX are important for all maturities, but inflation uncertainty is important only for medium and long maturities (five years and up).

3.3 Sensitivity Analysis

The results reported so far are based on quarterly data (since the SPF is a quarterly survey). As a sensitivity analysis, this section takes a look at monthly data. The inflation expectations here are from the Michigan survey and are for five-year inflation (and therefore combined with five-year interest rates). The Michigan survey has no direct information about inflation uncertainty, but it turns out that the implied volatility of bond options provides a proxy. (Recall that in table 1, the implied volatility is dominated by the inflation uncertainty.)

The basic specification is therefore to regress the five-year break-even deviation on the implied bond volatility and the off-the-run premium—and then to consider various extensions. Table 3 shows the results. Column 1 suggests that the basic specification does a decent job: the implied bond volatility has a significantly positive coefficient and might therefore capture the inflation risk premium, while the off-the-run premium has a significantly negative coefficient (as before).

Columns 2 and 3 add different proxies for the real liquidity premium: both the TIPS trading volume and the VIX have the expected signs and are significant. Columns 4 and 5 instead add alternative proxies for inflation risk. In column 5, the disagreement among the Michigan survey participants is (again) not significant. However, column 4 shows that a fitted monthly series of inflation uncertainty is. This series is created by first regressing the (quarterly) SPF series on inflation uncertainty on the (monthly) implied volatility and inflation disagreement—and then filling in the missing values with the fitted values.

Overall, the sensitivity analysis indicates that the quarterly and monthly frequencies give fairly similar results—although the lack of a clean measure of inflation uncertainty is a drawback of the monthly data.

4. Break-Even Inflation in Other Economies

Break-even inflation is used in many countries—and some even have extensive survey data on inflation uncertainty. This section briefly discusses that data for the United Kingdom and the euro area.

Table 3. Regression Results, Break-Even Deviation for the Five-Year Maturity, Monthly Data

	(1)	(2)	(3)	(4)	(5)
Inflation Uncertainty (Fitted)				2.37 (2.54)	
Implied Volatility	3.17 (2.30)	0.84 (0.67)	5.60 (3.37)	2.93 (2.25)	3.66 (2.60)
Inflation Disagreement					−0.59 (−1.69)
Off-the-Run Premium	−3.41 (−6.80)	−1.93 (−3.28)	−1.55 (−2.63)	−3.25 (−6.95)	−3.59 (−7.02)
TIPS Trade		4.70 (4.32)			
VIX			−0.03 (−4.98)		
Constant	−0.76 (−21.42)	−0.76 (−22.53)	−0.77 (−25.12)	−0.76 (−22.69)	−0.77 (−21.79)
R^2	0.32	0.44	0.50	0.36	0.35
Observations	103.00	103.00	103.00	103.00	103.00

Notes: This table shows regression coefficients and *t*-statistics (in parentheses) for monthly data 1997:M1–2008:M6. The dependent variable is the nominal interest rate minus the real interest rate and minus expected inflation (five-year horizon). All regressors (except the constant) are demeaned. The *t*-statistics are based on a Newey-West estimator with one lag.

The left panel of figure 5 shows UK break-even inflation (four years) and inflation expectations (two years), while the right panel shows inflation uncertainty and disagreement (based on the Bank of England's [BoE's] survey of external forecasters, also two-year horizon) from Boero, Smith, and Wallis (2008).[10] The survey data has a shorter horizon than the break-even inflation since the BoE only

[10]The break-even inflation rates are from the Bank of England's home page, www.bankofengland.co.uk/statistics/yieldcurve/archive.htm. The point forecasts are based on the average probabilities published in the Inflation Report (various issues). The data on uncertainty and disagreement was kindly provided by Kenneth F. Wallis. (The Bank of England does not yet publish the individual data from their survey.)

Figure 5. Data for UK Inflation: Break-Even Inflation, Inflation Expectations, and a Measure of Inflation Uncertainty

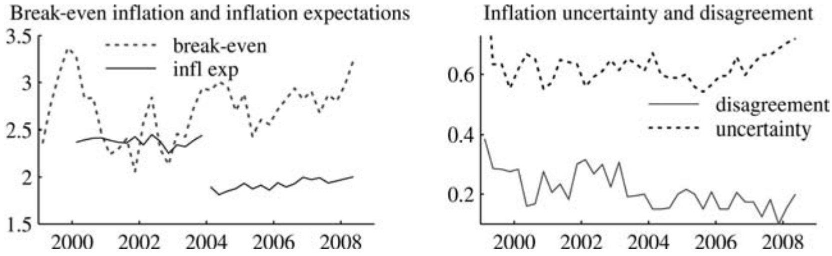

Notes: The left subfigure shows the four-year break-even inflation and the two-year inflation expectations (from the BoE survey of external forecasters). The right subfigure shows the inflation uncertainty and disagreement estimated from the individual answers (in the BoE survey of external forecasters) for the two-year horizon. The BoE asked about RPIX until November 2003 and CPI after that. The break-even inflation refers to RPI.

started asking about somewhat longer horizons in 2006. In addition, there is a break in the inflation expectations since the survey switched in early 2004 from asking about RPIX to CPI. In spite of these issues, there are interesting features of the data.

First, break-even inflation appears to vary much more than inflation expectations. This suggests that inflation risk and liquidity premia may have played an important role. For instance, break-even inflation moves away from inflation expectations—around the same time as uncertainty and disagreement are high. Second, both break-even inflation and inflation uncertainty seem to be somewhat more stable in the United Kingdom than in the United States (compare figures 2 and 4)—although the levels cannot be compared, as they refer to different horizons. This gives some additional support to Gürkaynak, Levin, and Swanson (2010), who argue that inflation expectations are more firmly anchored in the (inflation-targeting) United Kingdom than in the United States. (However, UK inflation uncertainty was distinctly higher before 2000—and the 1999:Q1 uncertainty is actually outside the chart.) In any case, this data deserves to be analyzed further.

Figure 6 shows similar data for the euro area. The break-even inflation rates are from inflation swaps from Beechey, Johannsen, and Levin (2010), and the other data is based on the individual

Figure 6. Data for Euro Inflation: Break-Even Inflation, Inflation Expectations, and Measures of Inflation Uncertainty and Disagreement

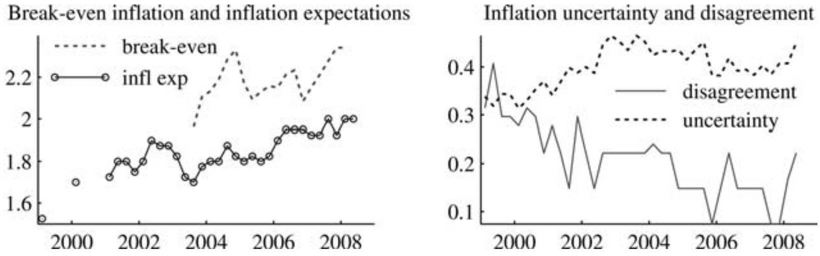

Notes: The left subfigure shows the break-even inflation (from inflation swaps) and the corresponding inflation expectations (from the ECB SPF) for a four-year period starting one year ahead. The right subfigure shows the inflation uncertainty and disagreement estimated from the individual answers (in the ECB SPF) for the two-year horizon.

answers to the ECB Survey of Professional Forecasters.[11] As for the United Kingdom and United States, the break-even inflation rates are for a longer horizon than the survey data.

Once again, break-even inflation moves more than inflation expectations—and there are some indications that the deviation is related to inflation uncertainty (although the sample is short). It also seems as if the euro-area break-even inflation rates and inflation uncertainty are somewhat more stable than in the United States.

Some work has been done on using the survey data in explaining the euro-area break-even inflation rates. In particular, García and Werner (2010) show, in an estimated no-arbitrage yield-curve model, that it is perhaps not inflation uncertainty, but the perceived asymmetry (skewness) of inflation risks, that affects inflation risk premia.

5. Concluding Remarks

This paper tries to explain the observed "break-even deviation" (the difference between break-even inflation and survey data on inflation expectations) by the off-the-run liquidity premium and the

[11] The data is available at www.aeaweb.org/forthcoming/output/accepted_MAC.php and www.ecb.int/stats/prices/indic/forecast/html/index.en.html.

survey data (from the Survey of Professional Forecasters) on inflation uncertainty. It is found that these variables are statistically and economically significant. Inflation uncertainty seems to capture an inflation risk premium, while the off-the-run liquidity premium proxies the real liquidity premium. This leads to important adjustments of break-even inflation.

The major comparative advantage of break-even inflation (relative to inflation surveys) is being a fast source of information. In particular, it helps analysts and policymakers to quickly gauge inflation expectations after important events. The findings in this paper suggest ways to adjust break-even inflation to get a better proxy of inflation expectations. In particular, if the event is associated with financial market turbulence, then this is likely to raise the real liquidity premium—which tends to decrease break-even inflation. The regression results presented above suggest how to adjust break-even inflation upward. In contrast, if the event is likely to increase inflation uncertainty, then this will drive up nominal rates—increasing break-even inflation. The regression results then show how to adjust break-even inflation down.

The main database of the paper—the Survey of Professional Forecasters—is only quarterly. Therefore, I also studied the monthly Reuters/Michigan Survey of Consumers, which contains inflation expectations but no direct measure of inflation uncertainty. The results suggest, however, that information from bond option prices may provide a reasonable proxy for inflation uncertainty, while data on disagreement (among the survey participants) about future inflation forecasts does not. There is also some evidence that the TIPS trading volume and the VIX index may help explain the liquidity premium of real bonds.

I hope that this paper contributes to the understanding of the economic factors behind the inflation risk and liquidity premia. It could also be of help to central bank staff and others who need to adjust break-even inflation. Further research might be able to find better high-frequency proxies for inflation uncertainty and thereby further improve current practices.

References

Beechey, M. J., B. K. Johannsen, and A. T. Levin. 2010. "Are Long-Run Inflation Expectations Anchored More Firmly in the

Euro Area than in the United States?" Forthcoming in *American Economic Journal: Macroeconomics*.

Boero, G., J. Smith, and K. F. Wallis. 2008. "Uncertainty and Disagreement in Economic Prediction: The Bank of England Survey of External Forecasters." *Economic Journal* 118 (530): 1107–27.

Carlstrom, C. T., and T. S. Fuerst. 2004. "Expected Inflation and TIPS." *Economic Commentary* (Federal Reserve Bank of Cleveland) (November): 1–4.

Chen, R.-R., B. Liu, and X. Cheng. 2005. "Inflation, Fisher Equation, and the Term Structure of Inflation Risk Premia: Theory and Evidence from TIPS." SSRN eLibrary.

Ciccarelli, M., and J. A. García. 2009. "What Drives Euro Area Break-Even Inflation Rates?" ECB Working Paper No. 996.

D'Amico, S., D. H. Kim, and M. Wei. 2008. "Tips from TIPS: The Informational Content of Treasury Inflation-Protected Security Prices." BIS Working Paper No. 248.

Federal Reserve Bank of Cleveland. 2010. "TIPS Expected Inflation Estimates." Available at http://www.clevelandfed.org/research/data/tips/index.cfm.

Federal Reserve Bank of New York. 2010. "Primary Dealers." Available at http://www.newyorkfed.org/markets/primarydealers.html.

Federal Reserve Bank of Philadelphia. 2010. "Survey of Professional Forecasters." Available at http://www.philadelphiafed.org/econ/spf/index.cfm.

Fleming, M. J., and N. Krishnan. 2009. "The Microstructure of the TIPS Market." Staff Report No. 414, Federal Reserve Bank of New York.

García, J. A., and A. Manzanares. 2007. "What Can Probability Forecasts Tell Us About Inflation Rules?" ECB Working Paper No. 825.

García, J. A., and T. Werner. 2010. "Inflation Risks and Inflation Risk Premia." ECB Working Paper No. 1162.

Giordani, P., and P. Söderlind. 2003. "Inflation Forecast Uncertainty." *European Economic Review* 47 (6): 1037–59.

Graveline, J. J., and M. R. McBrady. 2006. "Who Makes On-the-Run Treasuries Special?" Mimeo, Stanford Graduate School of Business.

Grishchenko, O. V., and J. Huang. 2008. "Inflation Risk Premium: Evidence from the TIPS Market." Mimeo, Penn State University.

Gürkaynak, R. S., A. Levin, and E. Swanson. 2010. "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the U.S., U.K., and Sweden." *Journal of the European Economic Association* 8 (6): 1208–42.

Gürkaynak, R. S., B. Sack, and J. H. Wright. 2008. "The TIPS Yield Curve and Inflation Compensation." Finance and Economics Discussion Series Paper No. 2008-05, Board of Governors of the Federal Reserve System.

Joyce, M., P. Lildholdt, and S. Sorensen. 2009. "Extracting Inflation Expectations and Inflation Risk Premia from the Term Structure: A Joint Model of the UK Nominal and Real Yield Curves." Working Paper No. 360, Bank of England.

Kajuth, F., and S. Watzka. 2008. "Inflation Expectations from Index-Linked Bonds: Correcting for Liquidity and Inflation Risk Premia." Mimeo, Ludwig-Maximilans-Universität Munich.

Lahiri, K., C. Teigland, and M. Zaporowski. 1988. "Interest Rates and the Subjective Probability Distribution of Inflation Forecasts." *Journal of Money, Credit, and Banking* 20 (2): 233–48.

McCulloch, J. H. 2008. "The US Real Term Structure of Interest Rates." Available at http://economics.sbs.ohio-state.edu/jhm/ts/ts.html.

Reschreiter, A. 2004. "Conditional Funding Costs of Inflation-Indexed and Conventional Government Bonds." *Journal of Banking and Finance* 28 (6): 1299–1318.

Reuters/University of Michigan. 2010. "Surveys of Consumers." Available at http://www.sca.isr.umich.edu/.

Roush, J. E. 2008. "The 'Growing Pains' of TIPS Issuance." Finance and Economics Discussion Series Paper No. 2008-08, Board of Governors of the Federal Reserve System.

Shen, P. 2006. "Liquidity Risk Premia and Breakeven Inflation Rates." *Economic Review* (Federal Reserve Bank of Kansas) (Second Quarter): 29–54.

U.S. Treasury. 2010. "Monthly Statement of Public Debt." Available at http://www.treasurydirect.gov/govt/reports/pd/mspd/mspd.htm.

Wright, J. H. 2008. "Term Premiums and Inflation Uncertainty: Empirical Evidence from an International Panel Dataset." Finance and Economics Discussion Series Paper No. 2008-25, Board of Governors of the Federal Reserve System.

Interest Rate Forecasts: A Pathology*

Charles A. E. Goodhart and Wen Bin Lim
Financial Markets Group
London School of Economics

This paper examines how well forecasters can predict the future time path of (policy-determined) short-term interest rates. Most prior work has been done using U.S. data; in this exercise we use forecasts made for New Zealand by the Reserve Bank of New Zealand (RBNZ) and those derived from money market yield curves in the United Kingdom. We broadly replicate recent U.S. findings for New Zealand and the United Kingdom, to show that such forecasts in New Zealand and the United Kingdom have been excellent for the immediate forthcoming quarter, reasonable for the next quarter, and useless thereafter. Moreover, when ex post errors are assessed depending on whether interest rates have been in an upward, or downward, section of the cycle, they are shown to have been biased and, apparently, inefficient. We attempt to explain those findings, and examine whether the apparent ex post forecast inefficiencies may still be consistent with ex ante forecast efficiency. We conclude, first, that the best forecast may be a hybrid containing a specific forecast for the next six months and a "no-change" assumption thereafter, and, second, that the modal forecast for interest rates, and maybe for other variables as well, is skewed, generally underestimating the likely continuation of the current phase of the cycle.

JEL Codes: C53, E17, E43, E47.

1. Introduction

The short-term policy interest rate has generally been adjusted in most developed countries, at least during the last twenty years or so, in a series of small steps in the same direction, followed by a pause

*Author contact: C.A.E. Goodhart, Financial Markets Group, Room R414, London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom. E-mail: c.a.goodhart@lse.ac.uk.

Figure 1. Official Cash Rate: Reserve Bank of New Zealand

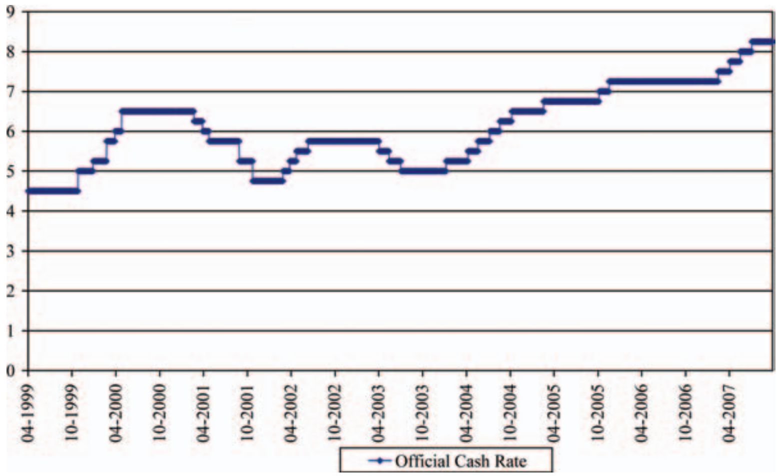

Source: Reserve Bank of New Zealand.

and then a, roughly, similar series of steps in the opposite direction. Figures 1 and 2 show the time path of policy rates for New Zealand and the United Kingdom, respectively.

On the face of it, such a behavioral pattern would appear quite easy to predict. Moreover, central bank behavior has typically been modeled by fitting a Taylor reaction function incorporating a lagged dependent variable with a large (often around 0.8 at a quarterly periodicity) and highly significant coefficient. But if this was, indeed, the reason for such gradualism, then the series of small steps should be highly predictable in advance.

The problem is that the evidence shows that they are *not* well predicted, beyond the next few months. There is a large body of, mainly American, literature to this effect, with the prime exponent being Glenn Rudebusch with a variety of co-authors; see in particular Rudebusch (1995, 2002, and 2006). Indeed, prior to the mid-1990s, there is some evidence that the market could hardly predict the likely path, or direction of movement, of policy rates over the next few months in the United States (see Rudebusch 1995 and 2002 and the literature cited there). More recently, with central banks having become much more transparent about their thinking, their

Figure 2. Official Bank Rate: Bank of England

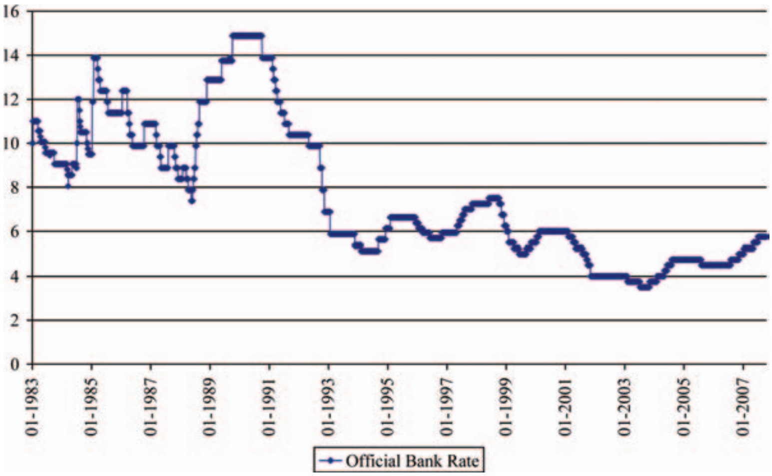

Source: Bank of England web site.

plans, and their intentions, market forecasts of the future path of policy rates have become quite good over the immediately forthcoming quarter, and better than a random-walk (no-change) assumption over the following quarter. But thereafter they remain as bad as ever (see Lange, Sack, and Whitesell 2003 and Rudebusch 2006).

We contribute to this literature first by extending the empirical analysis to New Zealand and the United Kingdom, though some similar work on UK data has already been done by Lildholdt and Wetherilt (2004). The work on New Zealand is particularly interesting, since the forecasts are *not* those derived from the money market but those made available by the Reserve Bank of New Zealand in their Monetary Policy Statements about their current expectations for their own future policies.

One of the issues relating to the question of whether a central bank should attempt to decide upon, and then publish, a prospective future path for its own policy rate, as contrasted with relying on the expected path implicit in the money market yield curve, is the relative precision of the two sets of forecasts. A discussion of the general issues involved is provided by Goodhart (2009). For an analytical discussion of the effects of the relative forecasting precision on that decision, see Morris and Shin (2002) and Svensson

(2006). An assessment of the effects of publicly announcing the forecast on market rates is given in Andersson and Hofmann (2009) and in Ferrero and Nobili (2009).

The question of the likely precision of a central bank's forecast of its own short-run policy rate is, however, at least in some large part, empirical. The Reserve Bank of New Zealand (RBNZ), a serial innovator in so many aspects of central banking, including inflation targeting and the transparency (plus sanctions) approach to bank regulation, was, once again, the first to provide a forecast of the (conditional) path of its own future policy rates. It began to do so in 2000:Q1. That gives twenty-eight observations between that date and 2006:Q4, our sample period. While still short, this is now long enough to undertake some preliminary tests to examine forecast precision.

Partly for the sake of comparison,[1] we also explore the accuracy of the implicit market forecasts of the path of future short-term interest rates in the United Kingdom. We use estimates provided by the Bank of England over the period 1992:Q4 until 2004:Q4. There are two such series, one derived from the London Interbank Offered Rate (LIBOR) yield curve and one from short-dated government debt. We base our choice between these on the relative accuracy of their forecasts. On this basis, as described in section 3, we chose, and subsequently used, the government debt series and its implied forecasts.

In the next section, section 2, we report and describe our data series. Then in section 3 of this paper we examine the predictive accuracy of these sets of interest rate forecasts. The results are closely in accord with the earlier findings in the United States.

[1]The United Kingdom and New Zealand (NZ) are different economies, and so one is *not* strictly comparing like with like. If one was, however, to compare the NZ implicit market forecast accuracy with that of the RBNZ forecast over the same period (a comparison which we hope that the RBNZ will do), the former will obviously be affected by the latter (and possibly vice versa). Again, if a researcher was to compare the implied accuracy of the market forecast *prior* to the introduction of the official forecast with the accuracy of the market/official forecast *after* the RBNZ had started to publish (another exercise that we hope that the RBNZ will undertake), then the NZ economy, their financial system, and the economic context may have changed over time. So one can *never* compare an implicit market forecast with an official forecast for interest rates on an exactly like-for-like basis. Be that as it may, we view the comparison of the RBNZ and the implied UK interest rate forecasts as illustrative, and not definitive in any way.

Figure 3. RBNZ Interest Rate Forecast (Ninety Days, Annualized Rate) Published in Successive Monetary Policy Statements

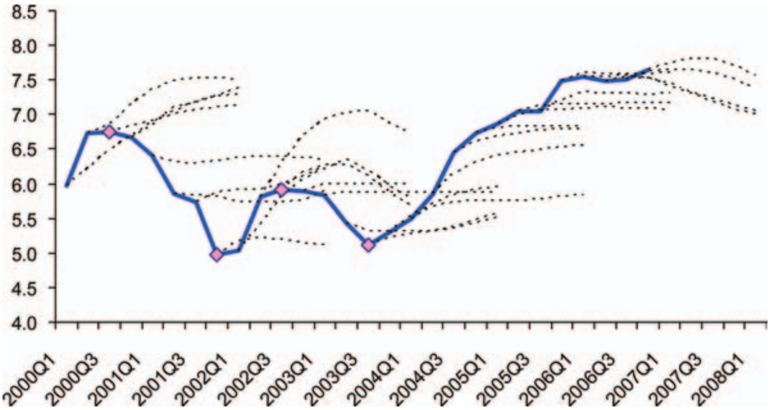

Notes: Turning points are marked by a diamond. The dating of these is discussed further in section 3.

Whether the forecast comes from the central bank or from the market, the predictive ability is good, by most econometric standards, over the first quarter following the date of the forecast; it is poor, but significantly better than a no-change, random-walk forecast, over the second quarter (from end-month 3 to end-month 6), and effectively useless from that horizon onward.

Worse, however, is to come. The forecasts, once beyond the end of the first quarter, are not only without value, they are, when compared with ex post outcomes, also strongly and significantly biased. This does not, however, necessarily mean that the forecasts were ex ante inefficient. We shall demonstrate in section 5 how ex post bias can yet be consistent with ex ante efficiency in forecasting.

This bias can actually be seen clearly in a visual representation of the forecasts. The RBNZ forecasts and outcome are shown in figure 3, and the UK forecast derived from the short-dated government debt yield curve and outcome is shown in figure 4.

What is apparent by simple inspection is that when interest rates are on an upward (downward) cyclical path, the forecast underestimates (overestimates) the actual subsequent path of interest rates. Much the same pattern is also observable in the United States (see

Figure 4. UK Interest Rate Forecast (Ninety Days, Annualized Rate) Derived from the Short-Dated Government Debt Yield Curve

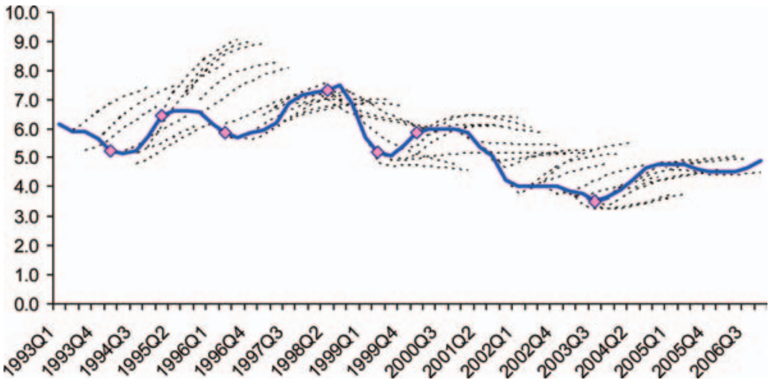

Rudebusch 2007) and Sweden (see Adolfson et al. 2007). One of the reasons why this bias has not been more widely recognized up till now is that the biases during up and down cyclical periods are almost exactly offsetting, so if an econometrician applies his or her tests to the complete time series (as usual) (s)he will find no aggregate sign of bias. The distinction between the bias in "up" and "down" periods is crucial. A problem with some time series—e.g., those for inflation—is that the division of the sample into "up," "down," and in some cases "flat" periods is not always easy, nor self-evident. But this is less so for short-term interest rates where the ex post timing of turning points is relatively easier.

The sequencing of this paper proceeds as follows. We report our database in section 2. We examine the accuracy of the interest rate forecasts in section 3. We continue in section 4 by assessing whether forecasts which appear ex post biased can still be ex ante efficient. Section 5 concludes.

2. The Database for Interest Rates

Our focus in this paper concerns the accuracy of forecasts for short-term policy-determined interest rates, measured in terms of unbiasedness and the magnitude of forecast error. We examine the data for two countries. We do so first for New Zealand, because this is the country with the longest available published series of official projections, as presented by the RBNZ in their quarterly Monetary Policy Statement. Our second country is the United Kingdom. In this case the Bank of England assumed unchanged future interests, from their current level, as the basis of their forecasts, until they moved onto a market-based estimate of future policy rates in November 2004. As described below, we considered the use of two alternative estimates of future (forecast) policy rates.

In New Zealand, policy announcements, and the release of projections, are usually made early in the final month of the calendar quarter, though the research work and discussions in their Monetary Policy Committee (MPC) will have mostly taken place a couple of weeks previously. Thus the Statement contains a forecast for inflation for the current quarter ($h = 0$), though that will have been made with knowledge of the outturn for the first month and some partial evidence for the second. The Policy Targets Agreement between the Treasurer and the Governor is specified in terms of the CPI, and the forecast is made in terms of the CPI. This does not, however, mean that the RBNZ focuses exclusively on the overall CPI in its assessment of inflationary pressures.

In New Zealand, the policy-determined rate is taken to be the ninety-day (three-month) rate, and the forecasts are for that rate. Thus the current-quarter interest rate observation contains nearly two months of actual ninety-day rates and just over one month of market forward one-month rates. If the MPC meeting results in a (revisable) decision to change interest rates in a way that is inconsistent with the prediction that was previously embedded in market forward interest rates, then the assumption for the current quarter can be revised to make the overall ninety-day track look consistent with the policy message. Finally, the policy interest rate can be adjusted, after the forecast is effectively completed, right up to the day before the Monetary Policy Statement; this was done in September 2001 after the terrorist attack. So, the interest rate forecast for the current quarter ($h = 0$) also contains a small extent of uncertain forecast.

The data for published official forecasts of the policy rate start in 2000:Q1. We show those data, the forecasts, and the resulting errors, for the policy rate in the appendix, tables 8 and 9. The data are shown in a format where the forecasts are shown in the same row as the actual to be forecast, so the forecast errors can be read off directly.

The British case is somewhat more complicated. In the past, during the years of our sample, the MPC used a constant forward forecast of the repo rate as the conditioning assumption for its forecasting exercise. Whether members of the MPC made any mental reservations about the forecast on account of a different subjective view about the future path of policy rates is an individual question that only they can answer personally. But it is hard to treat that constant path as a pure, most likely, forecast. At the same time, there are at least two alternative time series of implied market forecasts for future policy rates that are derived from the yield curve of short-dated government debt and from LIBOR. There are some complicated technical issues in extracting implied forecasts from market yield curves, and such yield curves can be distorted, especially the LIBOR yield curve, as experience since 2007 has clearly demonstrated. These problems relate largely to risk premia, notably credit and default risk; see Ferrero and Nobili (2009). The yield curve for government debt is (or rather has been) largely immune to such credit (default) risk, though it can be exposed to other risks, e.g., interest rate and liquidity risks.

We do not rehearse these difficulties here; instead we simply took these data from the Bank of England web site (see www.bankofengland.co.uk). For more information on the procedures used to obtain such implicit forecast series, see Anderson and Sleath (1999, 2001), Brooke, Cooper, and Scholtes (2000), and Joyce, Relleen, and Sorensen (2007). As will be reported in the next section, the government debt implicit market forecast series has had a more accurate forecast than the LIBOR series over our data period, 1992–2004, probably in part because the government series would not have incorporated a time-varying credit risk element; see Ferrero and Nobili (2009). Since the constant rate assumption was hardly a forecast, most of our work was done with the government debt implicit forecast series. This forecasts the three-month Treasury bill series. These series—actual, forecast, and errors (with the forecast lined up against the actual it was predicting)—are shown in the appendix, tables 10 and 11, for the government debt series (the other series for LIBOR is available from the authors on request).

3. How Accurate Are the Interest Rate Forecasts?

We began our examination of this question by running three regressions both for the NZ data series and for two sets of implied market forecasts for the United Kingdom, derived from the LIBOR and government debt yield curve, respectively. These regression equations were as follows:

$$IR(t+h) = C_1 + C_2 \text{ Forecast } (t, t+h) \tag{1}$$

$$IR(t+h) - IR(t) = C_1 + C_2 \ [\text{Forecast}(t, t+h) - IR(t)] \tag{2}$$

$$IR(t+h) - IR(t+h-1) = C_1 + C_2 \ [\text{Forecast}(t, t+h) - \text{Forecast}(t, t+h-1)], \tag{3}$$

where

$$IR(t) = \text{actual interest rate outturn at time } t$$

$$\text{Forecast}(t, t+h) = \text{forecast of } IR(t+h) \text{ made at time } t.$$

The first equation is essentially a Mincer-Zarnowitz regression (Mincer and Zarnowitz 1969) evaluating how well the forecast can predict the actual h-period-ahead interest rate outturn ($h = 0$ to n). If the forecast perfectly matches the actual interest rate outturn for every single period, we would expect to have $C_2 = 1$ and $C_1 = 0$. This can be seen as an evaluation of the bias of the forecast. Taking expectations on both sides, $E\{IR(t+h)\} = E\{C_1 + C_2[\text{Forecast}(t, t+h)\}$. A forecast is unbiased—i.e., $E\{IR(t+h)\} = E\{[\text{Forecast}(t, t+h)]\}$ for all t—if and only if $C_2 = 1$ and $C_1 = 0$. The second regression, by subtracting the interest rate level from both sides, allows us to focus our attention on the performance of the forecast interest rate difference $\{IR(t+h) - IR(t)\}$. It asks, as h increases, how accurately can the forecaster forecast h-quarter-ahead interest rate *changes* from the present level. The third regression is a slight twist on the second, focusing on one-period-ahead forecasts; the regression examines the forecast performance of one-period-ahead interest rate changes $\{IR(t+h) - IR(t+h-1)\}$ as h increases.

All three regressions assess the accuracy/biasness of interest rate forecasts from slightly different angles. An unbiased forecast necessarily implies a constant term of zero and a slope coefficient of one. We can test whether these conditions are fulfilled with a joint hypothesis test:

$$H_0\colon C_1 = 0 \text{ and } C_2 = 1.$$

With three equations, three data sets, and $h = 0$ to 5 for New Zealand and $h = 1$ to 8 for the UK series, we have some eighty-five regression results and statistical test scores to report.

We found that the regression results, estimated by OLS, for the implicit forecasts derived from the LIBOR yield curve were comprehensively worse than those from the government yield curve, or the RBNZ. These LIBOR results provided poor forecasts even for the first two quarters, and useless forecasts thereafter. There are several possible reasons for such worse forecasts—e.g., time-varying risk premia (Ferrero and Nobili 2009) or data errors in a short sample—but it is beyond the scope of this paper to try to track them down. These results can be found in Goodhart and Lim (2008) and, to save space, are not reported here. That reduces the number of regression results to sixteen in table 1 for the RBNZ and twenty-four in table 2 for the UK government yield curve.

These results show that the RBNZ forecast is excellent one quarter ahead but then becomes useless in forecasting the subsequent direction, or extent, of change. Thus the coefficient C_2 in equation (3) becomes -0.04 at $h = 2$ (with an R-squared of zero), and negative thereafter. When the equation is run in levels, rather than first differences—i.e., equation (1)—the excellent first-quarter forecast feeds through into a significantly positive forecast of the *level* in the next few quarters, though it is just the first-quarter forecast doing all the work. The Mincer-Zarnowitz test results[2] are also consistent with our findings. We failed to reject the joint hypothesis H_0 for up to a three-quarters-ahead forecast for equation (1) and up to a four-quarters-ahead forecast for equation (2). We reject H_0 for the quarters thereafter.

[2]These tests are reported in Goodhart and Lim (2008) but are omitted to save space here.

Table 1. Regression Results for New Zealand

$h =$	C_1 (p-value)	C_2 (p-value)	R-squared	DW
Equation (1)				
0	−0.01 (0.93)	1.00 (0.85)	0.99	1.77
1	−0.24 (0.64)	1.03 (0.74)	0.88	1.53
2	0.30 (0.75)	0.93 (0.63)	0.65	0.93
3	1.50 (0.25)	0.74 (0.19)	0.39	0.34
4	3.71 (0.03)	0.40 (0.02)	0.11	0.28
5	5.71 (0.00)	0.09 (0.00)	0.00	0.15
Equation (2)				
1	−0.16 (0.07)	1.61 (0.18)	0.35	1.61
2	−0.15 (0.31)	1.02 (0.95)	0.20	1.02
3	−0.09 (0.66)	0.73 (0.55)	0.10	0.45
4	0.13 (0.61)	0.11 (0.10)	0.00	0.47
5	0.37 (0.20)	−0.38 (0.01)	0.03	0.34
Equation (3)				
1	0.13 (0.07)	1.30 (0.33)	0.43	2.06
2	0.04 (0.65)	−0.04 (0.06)	0.00	1.24
3	0.07 (0.38)	−0.68 (0.04)	0.03	1.38
4	0.09 (0.28)	−1.29 (0.03)	0.07	1.37
5	0.09 (0.26)	−1.30 (0.02)	0.08	1.28

Note: The corresponding p-value is evaluated against the null hypothesis, H_0: $C_1 = 0$, $C_2 = 1$.

Table 2. UK Forecasts Derived from the Short-Term Government Yield Curve

$h =$	C_1 (p-value)	C_2 (p-value)	R-squared	DW
Equation (1)				
1	0.23 (0.25)	0.98 (0.64)	0.95	1.94
2	0.60 (0.07)	0.89 (0.06)	0.84	1.03
3	0.98 (0.03)	0.79 (0.01)	0.71	0.62
4	1.56 (0.00)	0.67 (0.00)	0.55	0.43
5	2.10 (0.00)	0.56 (0.00)	0.41	0.35
6	2.43 (0.00)	0.49 (0.00)	0.34	0.31
7	2.52 (0.00)	0.47 (0.00)	0.32	0.29
8	2.42 (0.00)	0.48 (0.00)	0.35	0.28
Equation (2)				
1	0.13 (0.02)	0.94 (0.70)	0.51	1.91
2	−0.01 (0.84)	0.86 (0.31)	0.50	1.04
3	−0.16 (0.09)	0.85 (0.25)	0.47	0.67
4	−0.28 (0.03)	0.73 (0.07)	0.36	0.48
5	−0.34 (0.03)	0.60 (0.01)	0.27	0.39
6	−0.37 (0.03)	0.51 (0.00)	0.22	0.35
7	−0.39 (0.02)	0.46 (0.00)	0.21	0.31
8	−0.43 (0.00)	0.46 (0.00)	0.24	0.28

(continued)

Table 2. (Continued)

$h =$	C_1 (p-value)	C_2 (p-value)	R-squared	DW
Equation (3)				
1	0.13 (0.02)	0.94 (0.70)	0.51	1.91
2	−0.13 (0.02)	0.87 (0.62)	0.25	1.19
3	−0.13 (0.04)	0.65 (0.14)	0.15	0.97
4	−0.09 (0.19)	0.43 (0.03)	0.05	0.83
5	−0.08 (0.21)	0.53 (0.14)	0.06	0.84
6	−0.08 (0.18)	0.73 (0.15)	0.08	0.80
7	−0.06 (0.34)	0.41 (0.58)	0.04	0.82
8	−0.05 (0.39)	0.58 (0.41)	0.03	0.76

Note: The corresponding p-value is evaluated against the null hypothesis, H_0: $C_1 = 0$, $C_2 = 1$.

Turning next to the United Kingdom implied forecasts from the government debt yield curve, what these tables indicate is that, in the first quarter after the forecast is made, the forecast precision of this derived forecast is mediocre (joint test for null hypothesis is rejected for $h = 3 - 8$), certainly significantly better than random walk (no change) but not nearly as good as the NZ forecast over its first quarter. However, this market-based forecast is also able to make a good forecast of the change in rates between Q1 and Q2 (whereas the RBNZ could not do that). The government yield forecast for $h = 2$ in table 2 is somewhat better than for $h = 1$. So the ability of the government yield forecast to predict the *level* of the policy rate two quarters (six months) hence is about the same or a little better than that of the RBNZ. Thereafter, from Q2 onward, the predictive ability of the government yield

Figure 5. Stylized Pattern of Relationships between Forecasts and Outturns of Macro Variables over the Cycle

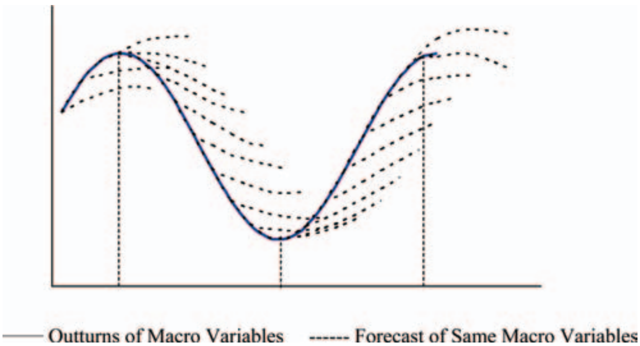

forecast becomes insignificantly different from zero, but at least the coefficients have the right sign (unlike the RBNZ).

The conclusion of this set of tests is that the precision of interest forecasts beyond the next quarter or two is approximately zero, whether they are made by the RBNZ or the UK market. Given the gradual adjustments in actual policy rates, this might seem surprising. Why does it happen? In order to answer this question, we start with a stylized fact. When one looks at most macroeconomic forecasts, and notably so for interest rates (see figures 3 and 4 above), they tend to follow a pattern. When the macro variable is rising, the forecast increasingly falls below it. When the macro variable is falling, the forecast increasingly lies above it. This pattern is shown again in illustrative form in figure 5.

So, if we divide the sample period into periods of rising and falling values for the variable of concern (in this case the interest rate), during up periods Actual minus Forecast will tend to be persistently positive, and during down periods Actual minus Forecast will tend to be persistently negative. There is, however, an important caveat. A forecast made during an up (down) period may extend several quarters beyond the turning point into the next down (up) period. Once a turning point has occurred, however, a forecast that was too high (low) during the continuing down (up) cycle can rapidly then become too low (high) once the cycle has switched direction. Clearly the tendency for Actual minus Forecast to be negative in an upturn will be most marked for forecasts made in an upturn so long as that upturn *continues*, i.e., until the next sign change from up to down, or vice versa. Nevertheless, we still expect on balance that forecasts made during an upturn (downturn) will tend to have positive (negative) Actual minus Forecast outturns even after such a sign change, but the result is clearly uncertain.[3] But the forecasts made for the policy rate in the next quarter (and to a lesser extent into the second quarter) are so good, especially for the next quarter for the RBNZ, that no such bias may exist.

As can be seen from figures 1 and 2, the official rate is frequently held constant for a period of a few months before there is a reversal of direction. So the exact date of reversal is somewhat uncertain. We chose a date during these months as the best alternative on the basis of other available contemporaneous evidence, notably the concurrent time path of market rates. But we also tested for robustness by taking the first and last dates of each flat period and rerunning the exercises. The latter made no difference; the results are available on request from the authors.

Perhaps the easiest way of demonstrating this result, suggested to us by Andrew Patton, is to run a regression of the forecast error, at various horizons, against two indicator variables, one for up periods (C_1) and one for down periods (C_2):[4]

$$[IR(t+h) - \text{Forecast}(t, t+h)] = C_1 \mathrm{I}_{up}(t+h) + C_2 \mathrm{I}_{down}(t+h), \tag{4}$$

where

$IR(t)$ = actual interest rate outturn at time t

$\text{Forecast}(t, t+h)$ = forecast of $IR(t+h)$ made at time t

$\mathrm{I}_{up}(t+h)$ is a dummy variable = 1 if time, $t+h$, is an "up" period; else 0

$\mathrm{I}_{down}(t+h)$ is a dummy variable = 1 if time, $t+h$, is a "down" period; else 0.

[3] When interest rates are volatile, and sign changes are more frequent, nothing useful can be said about the likely outcomes of Actual minus Forecast after a second sign change.

[4] In our original paper (Goodhart and Lim 2008), we did some additional and more complicated statistical exercises, looking at the number of errors of a particular sign, in "up" and "down" phases, their mean, standard deviation, and p-values. They are omitted here to save space.

Table 3. Results for New Zealand

A. Indicator Variable Is Based on State in NZ at Outturn Date (Whole Data Set)					
***H* =**	**Adj. R-sqr.**	**C_1**	**p-value**	**C_2**	**p-value**
Q1	0.41	0.06	0.26	−0.34	0.00
Q2	0.61	0.14	0.07	−0.69	0.00
Q3	0.58	0.23	0.06	−0.88	0.00
Q4	0.36	0.23	0.23	−0.99	0.00
Q5	0.27	0.24	0.33	−1.06	0.01
Q6	0.20	0.23	0.49	−1.07	0.05
Q7	0.03	0.13	0.79	−0.95	0.27
Q8	−0.30	0.04	0.97	−0.52	0.79
B. Indicator Variable Is Based on State in NZ at Outturn Date, but only Includes Period during Which Sign Is Unchanged					
***H* =**	**Adj. R-sqr.**	**C_1**	**p-value**	**C_2**	**p-value**
Q1	0.41	0.06	0.26	−0.34	0.00
Q2	0.76	0.22	0.00	−0.70	0.00
Q3	0.87	0.41	0.00	−1.13	0.00
Q4	0.81	0.56	0.00	−1.53	0.00
Q5	0.86	0.73	0.00	−2.13	0.00
Q6	—	—	—	—	—
Q7	—	—	—	—	—
Q8	—	—	—	—	—

Note: The corresponding p-value is evaluated against the null hypothesis, H_0: $C_1 = 0$, $C_2 = 0$.

The hypothesis is that the up-period indicator (C_1) is positive (actual > forecast) and the down-period indicator (C_2) is negative (actual < forecast).

The results for New Zealand are shown in table 3.

Turning next to the results for the UK government yield implied forecasts, we found similar results. In this case, however, the forecasts included some sizable *average* errors, whereby the forecasts implied that interest rates would tend to become higher than was the case in the historical event (actual < forecast). This average

Figure 6. Average UK Interest Rate Forecast Error

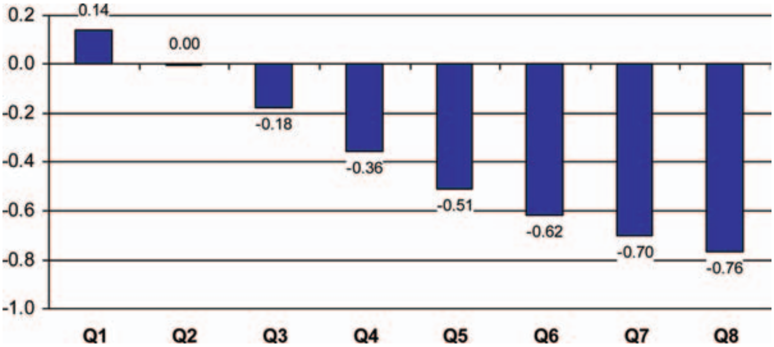

error tended to increase, approximately linearly, as the horizon (h) increased. This is shown in figure 6.

After correcting for this average error[5] and rerunning,[6] the results were as shown in table 4.

One of our referees kindly directed our attention to a related recent article, Ferrero and Nobili (2009). In this they regress excess returns (x), defined as forecast less actual ex post outcomes, for interest rates (futures) as a function of a business-cycle indicator (growth of output or employment expectations) and the current level of the futures rate, so that in their equation (4), p. 116,

$$x_{t+n}^{(n)} = a^{(n)} + \beta_n z_t + \gamma^n f_t^{(n)} + \varepsilon_{t+n}^{(n)},$$

where z_t is the business-cycle indicator, f_t is the level of the current futures rate, and β and γ are coefficients. In their table 2 (p. 118), table 5 (p. 127), and table 7 (p. 131), they find β to be negative, often significantly so, and γ to be usually significantly positive.

These authors cannot explain their own findings: “A theoretical analysis of the reasons behind the presence of forecast errors that are predictable and significantly countercyclical only in the United

[5]The tables using the unadjusted data—i.e., without correcting for the average error—are available on request from the authors.

[6]The average forecast error in New Zealand was much smaller and did not vary systematically with h. We ran similar adjusted regressions for New Zealand, but the results were closely similar to those shown in table 4.

Table 4. Results for United Kingdom, with Average Error Removed

A. Indicator Variable Is Based on State in UK at Outturn Date (Whole Data Set, with Average Forecast Error Removed)					
H **=**	**Adj. R-sqr.**	C_1	**p-value**	C_2	**p-value**
Q1	0.14	0.12	0.10	−0.08	0.16
Q2	0.25	0.26	0.01	−0.19	0.02
Q3	0.41	0.45	0.00	−0.38	0.00
Q4	0.22	0.41	0.02	−0.39	0.02
Q5	0.07	0.27	0.24	−0.30	0.17
Q6	0.01	0.04	0.89	−0.13	0.60
Q7	0.00	−0.19	0.50	−0.03	0.91
Q8	0.03	−0.40	0.16	0.01	0.97
B. Indicator Variable Is Based on State in UK at Outturn Date, but only Includes Period during Which Sign Is Unchanged, with Average Forecast Error Removed					
H **=**	**Adj. R-sqr.**	C_1	**p-value**	C_2	**p-value**
Q1	0.14	0.12	0.10	−0.08	0.16
Q2	0.32	0.28	0.01	−0.25	0.00
Q3	0.63	0.57	0.00	−0.55	0.00
Q4	0.70	0.79	0.00	−0.72	0.00
Q5	0.76	0.85	0.00	−0.88	0.00
Q6	0.81	0.76	0.00	−0.80	0.00
Q7	0.76	0.76	0.03	−0.89	0.00
Q8	0.67	0.52	0.23	−0.95	0.00

States lies beyond the scope of this paper" (Ferrero and Nobili 2009, p. 130). Our analysis here enables us to explain these findings; they are exactly what we would have expected given the ex post biases in forecasting over the cycle phases. As shown illustratively in figure 5, during the up (down) phases of the cycle, forecasts understate (overstate) ex post actuals systematically; hence β will be negative, though we too cannot explain why the euro zone exhibits less of this effect. Similarly, the expected futures rate will tend to be highest (lowest) at the top (bottom) of the cycle. As figure 5 again shows, this is when the forecast bias has forecast greater (less) than actual, so γ should be positive. The explanation of the Ferrero/Nobili results is, in our view, not due to time-varying risk premia, but to systematic ex post biases in the forecasting process over cycle phases. We are particularly grateful for having been given the chance to relate our work here to another strand in the literature.

What all these results show is as follows:

(i) The official and market forecasts of interest rates that we have studied here have significant predictive power over the next two quarters, but virtually none thereafter. When forecast precision is effectively zero, as after two quarters hence, it is perhaps best to acknowledge this, e.g., by the central bank using either a "no-change" thereafter assumption, or the implied market forecast, for the more distant forecasts.[7]

(ii) These interest rate forecasts are systematically biased, underestimating future policy rates during upturns and overestimating them during downturns. We shall now proceed to explore reasons why this might have been so in sections 4 and 5.

4. Can One Forecast the Forecasters?

In the preceding sections, we have shown that interest rate forecasts in the United Kingdom and New Zealand during this time period systematically underpredicted the time series during cyclical phases of upward movement, and similarly overpredicted during downswings.

In this section we seek to address the question of why these (most?) forecasts exhibit this tendency.[8] The answer that we propose is that (most) macroeconomic variables are expected (by most

[7]The choice may depend on the confidence with which the official forecasters hold their longer-dated forecasts. There is, however, a danger that the official forecasters have excessive confidence in their own forecasting abilities *and* that private-sector forecasters likewise place excessive weight on such official forecasts (Morris and Shin 2002). However, the finding by Ferrero and Secchi (2009) that long-term expectations on future interest rates react significantly only to short-term central bank interest rate forecasts, and not to their longer-term projections, suggests that market agents may well realize that such longer-term projections rarely contain any valuable information.

[8]In our original work we extended our research to cover inflation forecasts as well. These also exhibited the same syndrome. In order to save space and to enhance focus, we have, however, omitted those results from this paper. A more extended version of this paper, which explores not only the (errors in the) inflation forecasts in New Zealand and the United Kingdom but also the relationships between the errors in the inflation forecasts and those in the interest rate forecasts, is given in Goodhart and Lim (2009).

economists[9]) to revert to some longer-term equilibrium, ceteris paribus. Indeed, it is hard to see how forecasting could be done in the absence of a concept of (long-run) equilibrium. But at any particular point of time, macroeconomic variables will be subject to momentum, whose current force is quite difficult to assess accurately and which will be subject to unforeseeable future shocks. Thus we posit that these (most) forecasts will be subject to two main elements, an autoregressive component and a mean-reverting (back to equilibrium) component. Such a combination is bound to give us the general pattern that we have found in practice. So long as the phase remains upward (downward), the mean-reverting element in the forecast will tend to pull the forecast below (above) the actual track of the variable, but, of course, as the eventual turning point draws closer, it will predict far better than a pure autoregressive forecast.

During the periods under examination, an inflation-targeting regime was in operation in both New Zealand and the United Kingdom, so the equilibrium to which the inflation rate would revert would have been close to target, and about $2\frac{1}{2}$ percent above that for the nominal interest rate, assuming an equilibrium real interest rate of $2\frac{1}{2}$ percent. But for our purposes here, we do not assume to know what the equilibrium interest rate is, and have simply taken the arithmetic average of the study period as an estimation of the "mean-reverting" point.[10] The nature of the autoregressive process for each series, and the coefficients for combining the autoregressive and the mean-reverting components into an implied forecast are unknown and for determination. Initially we shall assume that the forecasters make an efficient, unbiased prediction of both factors. Thus we estimate for each series

$$IR(t+1) - IR(t) = B_1 * [IR(t) - IR(t-1)] + B_2 * [IR(t) - \overline{IR}], \quad (5)$$

[9]Not all economists have such expectations. A few, "heterodox," economists challenge whether equilibria necessarily exist, notably Paul Davidson and Basil Moore.

[10]We tested this by trying values of the mean-reverting point with +1 percent and −1 percent of the average, and it made no significant difference to the coefficients for B_1 and B_2, as well as for the results in table 6 and table 7. These results are available on request from the authors.

Table 5. Estimated Coefficients

Equation (1)	B_1			B_2			Regression Statistics		
	Coef.	*t*-stats	p-value	Coef.	*t*-stats	p-value	Adj. R-sqr.	SE	Obs.
UK Interest Rate	0.66	6.30	0.00	−0.09	−2.54	0.01	0.4175	0.2539	54
NZ Interest Rate	0.49	4.90	0.00	−0.13	−3.58	0.00	0.3403	0.6326	66

Figure 7. NZ Interest Rate: Comparison between Outturn, Actual and Implied Forecast

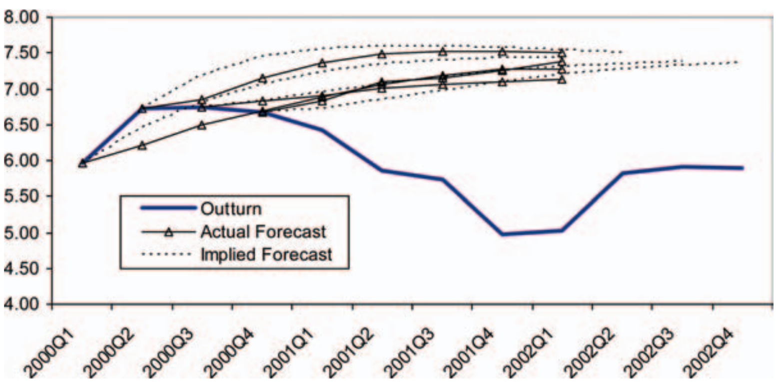

where

$$IR(t) = \text{actual interest rate outturn at time } t$$

$$\overline{IR} = \text{average interest rate outturn over the study period.}$$

The estimated coefficients are shown in table 5, where B_1 can be understood as the autoregressive coefficient, and B_2 as the mean-reversion coefficient.

Now we have a simplified model of how forecasts are done. The next step is to compare it with the actual forecasts. We do this first diagrammatically. For illustration, we have provided the diagrammatical comparison for the period between 2000:Q1 and 2002:Q4. The diagrams, figure 7 for New Zealand and figure 8 for the United

Figure 8. UK Interest Rate: Comparison between Outturn, Actual and Implied Forecast

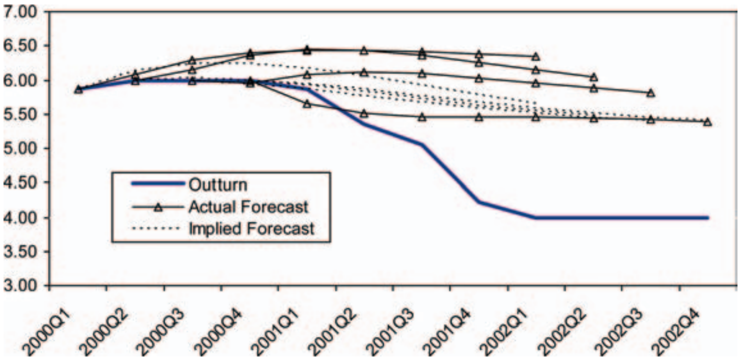

Kingdom, show quite a close relationship between the actual and our implied (from our simple model) forecast. Quarters beyond $t+1$ are estimated recursively.

We then evaluate the implied forecast *changes* against the actual forecast *changes* via regression analysis over the whole study period:

$$\begin{aligned}&\text{Actual Forecast } (t, t+h) - IR(t)\\&\quad = C_1 + C_2[\text{Implied Forecast } (t, t+h) - IR(t)\\&\quad\; i = 1 - 8.\end{aligned} \tag{6}$$

The hypothesis is that $C_1 = 0$ and $C_2 = 1$. The t-stats for C_2 in table 6 for New Zealand and table 7 for the United Kingdom relate to the coefficient's deviation from unity, *not* from zero.

But the regressions, and a closer inspection of the diagrams, indicated a systematic problem, separating the implied from the actual forecast. This was that the "true" coefficient of mean reversion during these years was greater than that used by the actual forecasters; i.e., the implied forecast flattened out near the equilibrium level faster than the actual forecasters expected. An indicative diagram for the six-quarters-ahead implied forecast for the UK interest rate showing this is given in figure 9.[11]

[11] Similar figures for NZ interest rates and for the UK series, both inflation and interest rates, are available in Goodhart and Lim (2009).

Table 6. NZ Interest Rate: Evaluation of Implied Forecast and Actual Forecast

	C_1			C_2			Regression Statistics		
h	**Coef.**	***t*-stats**	**p-value**	**Coef.**	***t*-stats**	**p-value**	**Adj. R-sqr.**	**SE**	**Obs.**
1	0.05	1.76	0.09	0.48	−4.99	0.00	0.44	0.10	28
2	0.01	0.14	0.89	0.57	−3.87	0.00	0.48	0.17	28
3	−0.03	−0.41	0.69	0.52	−4.20	0.00	0.43	0.24	28
4	−0.04	−0.40	0.69	0.44	−4.95	0.00	0.35	0.30	28
5	−0.05	−0.44	0.66	0.40	−5.35	0.00	0.30	0.35	28
6	−0.13	−0.81	0.43	0.43	−4.36	0.00	0.33	0.38	21
7	−0.21	−1.02	0.33	0.48	−3.22	0.01	0.38	0.41	14
8	−0.31	−0.91	0.41	0.50	−2.12	0.09	0.36	0.48	7

Table 7. UK Interest Rate: Evaluation of Implied Forecast and Actual Forecast

	C_1			C_2			Regression Statistics		
h	**Coef.**	***t*-stats**	**p-value**	**Coef.**	***t*-stats**	**p-value**	**R-sqr.**	**SE**	**Obs.**
1	−0.16	−5.51	0.00	0.89	−0.93	0.36	0.65	0.17	34
2	−0.03	−0.49	0.63	0.81	−1.33	0.19	0.42	0.39	44
3	0.17	1.96	0.06	0.70	−1.84	0.07	0.28	0.59	47
4	0.31	2.82	0.01	0.59	−2.42	0.02	0.20	0.75	47
5	0.43	3.27	0.00	0.50	−2.92	0.01	0.14	0.89	47
6	0.51	3.52	0.00	0.44	−3.33	0.00	0.11	0.99	47
7	0.58	3.66	0.00	0.40	−3.64	0.00	0.09	1.07	47
8	0.63	3.74	0.00	0.37	−3.86	0.00	0.08	1.14	47

Incidentally, the implied forecasts often did better in predicting the outturns than the actual forecasts. The results are available from the authors on request. This is not, however, so surprising since the implied forecasts are obtained by finding the coefficients that best explained the ex post outturns, i.e., data mining. So we place no emphasis on this finding.

The actual forecasters placed less weight on mean reversion than appeared to be the case in our constructed implied forecasts. That

Figure 9. UK Interest Rate, Six-Quarters-Ahead Forecast

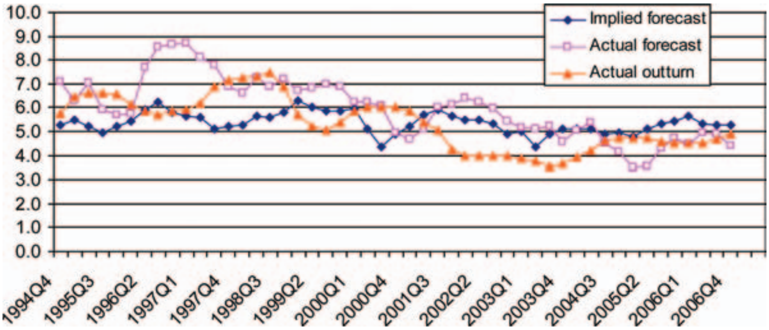

forecasters should have underestimated the speed of reversion to the mean is itself both plausible and understandable during these years. This was, after all, the period of the Great Moderation. A possible definition of such a Great Moderation is a period when the key macroeconomic time series revert to their (desired) equilibrium somewhat faster than in the past or than currently expected.

Most macro variables are cyclical, but, as any forecaster knows only too well, it is extraordinarily difficult to predict turning points. Hence a forecast which combines a weighted average of autoregressive continuation and mean reversion is likely to be optimal. It should minimize the likelihood of a really big error, and will be unbiased over the medium and longer run. So the behavior of the forecasters in seeking to estimate the likely mean outturn is, we would argue, appropriate.

Where our findings do indicate that there is a need for improvement is with the fan chart, or probability distribution, of future outcomes. This is usually shown as a symmetric single-peaked distribution, often akin to a normal distribution with mode, mean, and median at the same point.

Our results show that this will *generally not* be the case. The most probable outcome is that the cyclical phase will continue. Hence in an upturn (downturn), the most probable outcome is that (inflation and) interest rates will turn out to be systematically above (below) the mean forecast. But this is balanced by a smaller probability that the cycle will turn within this interval. But if there should be such a turning point, following an upturn (downturn) phase, then the forecasts will considerably overstate (understate) the subsequent downward (upward) movement.

5. Conclusions

In this paper we have demonstrated that, in the two countries and short data periods studied, the forecasts of interest rates had little or no informational value when the horizon exceeded two quarters (six months), though they were good in the next quarter and reasonable in the second quarter out. Moreover, all the forecasts were ex post and, systematically, inefficient, underestimating (overestimating) future outturns during up (down) cycle phases. The main reason for this is that forecasters cannot predict the timing of cyclical turning points, and hence predict future developments as a convex combination of autoregressive momentum and a reversion to equilibrium.

There are, perhaps, two main conclusions that can be drawn from this. The first is that official interest rate forecasts should probably be presented in hybrid form. MPCs and markets can make reasonable forecasts of interest rates up to two (at an extreme pinch, three) quarters hence. These should, indeed, be the basis of forecasts. Beyond that horizon, they are rarely able to do so, and that too should be acknowledged. Unless the authorities have a particular reason for exhibiting confidence in their own longer-dated forecasts, those same (longer-dated) forecasts should be presented in a specifically formulaic manner, e.g., constant or based on implied forward market rates.

The second conclusion is that the resulting interest (and inflation) forecast is generally *not* modal. It is biased, underestimating (overestimating) in upturns (downturns), because the forecaster is protecting himself or herself against extreme errors by assuming a (roughly constant) small probability of a turning point in the cycle occurring in each quarter. Consequently the *most likely* outturn in any expansionary phase is that output, inflation, and interest rates will turn out *above* forecast (vice versa in a downturn). The conclusion that we would draw from this is that policy needs to be normally somewhat more aggressive than the mean forecast would indicate (raising rates in booms, cutting rates in recessions), but that the policymakers need to be alert to (unpredictable) turning points and therefore to the occasional need to reverse course abruptly.

Appendix

Table 8. RBNZ Interest Rate Forecast

Date	Interest Rate	State[a]	$r(t,t)$	$r(t-1,t)$	$r(t-2,t)$	$r(t-3,t)$	$r(t-4,t)$	$r(t-5,t)$	$r(t-6,t)$	$r(t-7,t)$	$r(t-8,t)$
2000:Q1	5.97	1	5.86	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
2000:Q2	6.73	1	6.46	6.21	N/A	N/A	N/A	N/A	N/A	N/A	N/A
2000:Q3	6.74	1	6.83	6.84	6.49	N/A	N/A	N/A	N/A	N/A	N/A
2000:Q4	6.67	−1	6.64	6.83	7.15	6.70	N/A	N/A	N/A	N/A	N/A
2001:Q1	6.42	−1	6.50	6.84	6.91	7.36	6.88	N/A	N/A	N/A	N/A
2001:Q2	5.85	−1	5.84	6.31	7.10	7.01	7.48	7.05	N/A	N/A	N/A
2001:Q3	5.74	−1	5.79	5.83	6.30	7.16	7.07	7.53	7.19	N/A	N/A
2001:Q4	4.97	−1	5.07	5.87	5.81	6.34	7.26	7.10	7.53	7.27	N/A
2002:Q1	5.03	1	4.91	5.18	5.90	5.74	6.38	7.38	7.13	7.51	7.28
2002:Q2	5.82	1	5.72	5.41	5.22	5.92	5.74	6.39	N/A	N/A	N/A
2002:Q3	5.91	1	5.97	6.30	5.81	5.20	5.98	5.73	6.38	N/A	N/A
2002:Q4	5.90	−1	6.00	6.16	6.70	6.08	5.14	6.10	5.76	6.36	N/A
2003:Q1	5.83	−1	5.88	6.00	6.26	6.93	6.22	5.12	6.23	5.90	6.35
2003:Q2	5.44	−1	5.47	5.88	6.00	6.27	7.03	6.34	N/A	N/A	N/A
2003:Q3	5.12	−1	5.12	5.32	5.88	6.00	6.11	7.04	6.18	N/A	N/A
2003:Q4	5.29	1	5.32	5.22	5.31	5.88	6.00	5.88	6.87	5.96	N/A
2004:Q1	5.49	1	5.51	5.54	5.28	5.31	5.88	6.00	5.69	6.72	5.79
2004:Q2	5.86	1	5.76	5.67	5.71	5.31	5.32	5.88	N/A	N/A	N/A
2004:Q3	6.44	1	6.35	6.14	5.73	5.82	5.37	5.36	5.88	N/A	N/A

(continued)

Table 8. (Continued)

Date	Interest Rate	State[a]	$r(t,t)$	$r(t-1,t)$	$r(t-2,t)$	$r(t-3,t)$	$r(t-4,t)$	$r(t-5,t)$	$r(t-6,t)$	$r(t-7,t)$	$r(t-8,t)$
2004:Q4	6.73	1	6.74	6.61	6.31	5.75	5.90	5.47	5.44	5.88	N/A
2005:Q1	6.86	1	6.80	6.80	6.68	6.40	5.75	5.95	5.57	5.52	5.88
2005:Q2	7.04	1	7.00	7.00	6.83	6.73	6.45	5.75	N/A	N/A	N/A
2005:Q3	7.05	1	7.05	7.12	7.07	6.82	6.76	6.49	5.77	N/A	N/A
2005:Q4	7.49	1	7.47	7.21	7.15	7.07	6.83	6.78	6.53	5.81	N/A
2006:Q1	7.55	1	7.57	7.61	7.32	7.14	7.09	6.82	6.78	6.54	5.84
2006:Q2	7.48	1	7.49	7.55	7.59	7.31	7.15	7.10	N/A	N/A	N/A
2006:Q3	7.51	1	7.48	7.55	7.56	7.58	7.30	7.16	7.09	N/A	N/A
2006:Q4	7.64	1	7.62	7.62	7.53	7.53	7.59	7.29	7.17	7.09	N/A

[a] "+1" indicates an "up" period, i.e., a period of rising interest rate; "−1" indicates a "down" period, i.e., a period of declining interest rate.
Source: The NZ data are taken from their quarterly Monetary Policy Statements. For more detail, see section 2.

Table 9. RBNZ Interest Rate Forecast Error (Table Updated)

Forecast Error	$r(t,t)$	$r(t-1,t)$	$r(t-2,t)$	$r(t-3,t)$	$r(t-4,t)$	$r(t-5,t)$	$r(t-6,t)$	$r(t-7,t)$	$r(t-8,t)$
2000:Q1	0.11	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
2000:Q2	0.27	0.52	N/A	N/A	N/A	N/A	N/A	N/A	N/A
2000:Q3	−0.09	−0.10	0.25	N/A	N/A	N/A	N/A	N/A	N/A
2000:Q4	0.03	−0.16	−0.48	−0.03	N/A	N/A	N/A	N/A	N/A
2001:Q1	−0.08	−0.42	−0.49	−0.94	−0.47	N/A	N/A	N/A	N/A
2001:Q2	0.01	−0.46	−1.25	−1.16	−1.63	−1.20	N/A	N/A	N/A
2001:Q3	−0.05	−0.09	−0.56	−1.42	−1.33	−1.79	−1.46	N/A	N/A
2001:Q4	−0.10	−0.90	−0.84	−1.37	−2.29	−2.13	−2.56	−2.30	N/A
2002:Q1	0.12	−0.15	−0.87	−0.70	−1.34	−2.35	−2.10	−2.48	−2.25
2002:Q2	0.10	0.41	0.60	−0.10	0.08	−0.57	N/A	N/A	N/A
2002:Q3	−0.06	−0.38	0.10	0.72	−0.07	0.18	−0.47	N/A	N/A
2002:Q4	−0.10	−0.26	−0.80	−0.18	0.75	−0.21	0.14	−0.47	N/A
2003:Q1	−0.05	−0.17	−0.43	−1.11	−0.39	0.71	−0.41	−0.07	−0.52
2003:Q2	−0.03	−0.44	−0.56	−0.84	−1.59	−0.90	N/A	N/A	N/A
2003:Q3	0.00	−0.20	−0.76	−0.88	−0.98	−1.91	−1.06	N/A	N/A
2003:Q4	−0.03	0.07	−0.02	−0.59	−0.71	−0.59	−1.58	−0.67	N/A
2004:Q1	−0.02	−0.05	0.22	0.19	−0.39	−0.51	−0.20	−1.23	−0.29
2004:Q2	0.10	0.19	0.15	0.55	0.54	−0.02	N/A	N/A	N/A

(continued)

Table 9. (Continued)

Forecast Error	$r(t, t)$	$r(t-1, t)$	$r(t-2, t)$	$r(t-3, t)$	$r(t-4, t)$	$r(t-5, t)$	$r(t-6, t)$	$r(t-7, t)$	$r(t-8, t)$
2004:Q3	0.09	0.30	0.71	0.62	1.07	1.08	0.56	N/A	N/A
2004:Q4	−0.01	0.11	0.41	0.98	0.82	1.26	1.29	0.85	N/A
2005:Q1	0.06	0.06	0.18	0.46	1.11	0.91	1.29	1.34	0.98
2005:Q2	0.04	0.05	0.21	0.31	0.59	1.29	N/A	N/A	N/A
2005:Q3	0.00	−0.07	−0.02	0.23	0.29	0.56	1.28	N/A	N/A
2005:Q4	0.02	0.29	0.35	0.42	0.66	0.71	0.97	1.68	N/A
2006:Q1	−0.02	−0.06	0.23	0.41	0.46	0.73	0.77	1.01	1.71
2006:Q2	−0.01	−0.07	−0.11	0.17	0.32	0.38	N/A	N/A	N/A
2006:Q3	0.03	−0.04	−0.05	−0.07	0.21	0.35	0.42	N/A	N/A
2006:Q4	0.02	0.03	0.11	0.11	0.05	0.36	0.48	0.56	N/A

Note: Error is Actual minus Forecast.

Table 10. UK Interest Rate Forecast Implied by Government Yield Curve

Date	Interest Rate	State[a]	$r(t-1,t)$	$r(t-2,t)$	$r(t-3,t)$	$r(t-4,t)$	$r(t-5,t)$	$r(t-6,t)$	$r(t-7,t)$	$r(t-8,t)$
1993:Q1	6.13	−1	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
1993:Q2	5.88	−1	N/A	5.95	N/A	N/A	N/A	N/A	N/A	N/A
1993:Q3	5.88	−1	N/A	5.22	6.18	N/A	N/A	N/A	N/A	N/A
1993:Q4	5.66	−1	N/A	5.60	5.36	6.56	N/A	N/A	N/A	N/A
1994:Q1	5.22	−1	N/A	5.12	6.02	5.66	6.85	N/A	N/A	N/A
1994:Q2	5.13	1	N/A	5.14	5.17	6.43	5.98	7.07	N/A	N/A
1994:Q3	5.24	1	N/A	4.77	5.17	5.38	6.76	6.28	7.24	N/A
1994:Q4	5.75	1	N/A	5.36	4.94	5.30	5.65	7.03	6.56	7.40
1995:Q1	6.45	1	N/A	6.55	6.08	5.21	5.49	5.92	7.26	6.81
1995:Q2	6.63	−1	N/A	N/A	7.23	6.73	5.49	5.71	6.17	7.47
1995:Q3	6.63	−1	N/A	7.14	7.42	7.80	7.27	5.75	5.93	6.40
1995:Q4	6.58	−1	6.49	6.97	7.73	7.97	8.24	7.69	5.98	6.14
1996:Q1	6.13	−1	N/A	6.76	7.39	8.20	8.39	8.57	8.01	6.18
1996:Q2	5.87	−1	5.68	6.16	7.08	7.73	8.52	8.68	8.83	8.26
1996:Q3	5.69	1	N/A	5.64	6.29	7.39	7.95	8.72	8.88	9.02
1996:Q4	5.86	1	5.60	N/A	5.84	6.50	7.63	8.09	8.85	9.00
1997:Q1	5.94	1	N/A	5.74	6.42	6.12	6.71	7.82	8.18	8.93
1997:Q2	6.20	1	N/A	6.63	6.01	6.74	6.37	6.90	7.96	8.24
1997:Q3	6.87	1	6.22	N/A	6.88	6.34	7.01	6.60	7.06	8.06
1997:Q4	7.15	1	6.87	6.51	6.43	7.04	6.62	7.24	6.80	7.19
1998:Q1	7.25	1	7.26	6.95	6.67	6.57	7.13	6.86	7.43	6.97

(continued)

Table 10. (Continued)

Date	Interest Rate	State[a]	$r(t-1,t)$	$r(t-2,t)$	$r(t-3,t)$	$r(t-4,t)$	$r(t-5,t)$	$r(t-6,t)$	$r(t-7,t)$	$r(t-8,t)$
1998:Q2	7.33	1	7.00	7.22	6.99	6.76	6.66	7.19	7.04	7.58
1998:Q3	7.50	−1	6.94	6.69	7.10	6.99	6.80	6.73	7.23	7.19
1998:Q4	6.86	−1	7.21	6.71	6.51	7.00	6.98	6.83	6.78	7.25
1999:Q1	5.69	−1	6.10	6.96	6.53	6.39	6.93	6.98	6.85	6.82
1999:Q2	5.20	−1	4.80	5.79	6.69	6.41	6.30	6.87	6.98	6.87
1999:Q3	5.07	1	4.89	4.69	5.51	6.46	6.31	6.21	6.82	6.98
1999:Q4	5.40	1	4.89	4.89	4.71	5.28	6.26	6.21	6.13	6.77
2000:Q1	5.87	1	5.37	5.10	4.94	4.72	5.10	6.09	6.13	6.05
2000:Q2	6.00	−1	6.07	5.79	5.45	5.02	4.70	4.96	5.93	6.05
2000:Q3	6.00	−1	6.14	6.29	6.00	5.75	5.08	4.66	4.86	5.80
2000:Q4	6.00	−1	5.95	6.36	6.40	6.09	5.93	5.11	4.61	4.77
2001:Q1	5.86	−1	5.65	6.08	6.44	6.43	6.13	6.02	5.13	4.56
2001:Q2	5.36	−1	5.34	5.52	6.12	6.43	6.43	6.13	6.06	5.13
2001:Q3	5.05	−1	4.90	5.16	5.47	6.09	6.36	6.41	6.10	6.06
2001:Q4	4.23	−1	4.66	4.89	5.14	5.46	6.03	6.26	6.38	6.05
2002:Q1	4.00	−1	3.77	4.83	4.95	5.14	5.46	5.96	6.15	6.34
2002:Q2	4.00	−1	4.01	3.92	5.01	5.02	5.15	5.44	5.89	6.04
2002:Q3	4.00	−1	4.17	4.41	4.14	5.12	5.08	5.14	5.42	5.82

(continued)

Table 10. (Continued)

Date	Interest Rate	State[a]	$r(t-1,t)$	$r(t-2,t)$	$r(t-3,t)$	$r(t-4,t)$	$r(t-5,t)$	$r(t-6,t)$	$r(t-7,t)$	$r(t-8,t)$
2002:Q4	4.00	−1	3.74	4.59	4.68	4.32	5.19	5.12	5.13	5.39
2003:Q1	3.85	−1	3.72	3.80	4.90	4.85	4.47	5.22	5.14	5.12
2003:Q2	3.75	−1	3.38	3.68	3.98	5.12	4.97	4.59	5.24	5.15
2003:Q3	3.53	−1	3.34	3.27	3.76	4.19	5.27	5.04	4.69	5.24
2003:Q4	3.65	1	3.36	3.27	3.24	3.89	4.37	5.37	5.09	4.77
2004:Q1	3.91	1	3.96	3.60	3.28	3.29	4.02	4.52	5.44	5.12
2004:Q2	4.22	1	3.95	4.18	3.84	3.35	3.38	4.13	4.64	5.49
2004:Q3	4.65	1	4.42	4.10	4.35	4.03	3.44	3.49	4.24	4.73
2004:Q4	4.75	1	4.80	4.68	4.19	4.49	4.18	3.54	3.60	4.33

[a] "+1" indicates an "up" period, i.e., a period of rising interest rate; "−1" indicates a "down" period, i.e., a period of declining interest rate.
Source: The UK data are taken from the Bank of England web site, www.bankofengland.co.uk. For more detail, see section 2.

Table 11. Implied Error from UK Government Yield Forecast (Table Updated)

Forecast Error	$r(t-1,t)$	$r(t-2,t)$	$r(t-3,t)$	$r(t-4,t)$	$r(t-5,t)$	$r(t-6,t)$	$r(t-7,t)$	$r(t-8,t)$
1993:Q1	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
1993:Q2	N/A	N/A	−0.07	N/A	N/A	N/A	N/A	N/A
1993:Q3	N/A	N/A	0.66	−0.30	N/A	N/A	N/A	N/A
1993:Q4	N/A	N/A	0.06	0.30	−0.90	N/A	N/A	N/A
1994:Q1	N/A	N/A	0.10	−0.80	−0.44	−1.63	N/A	N/A
1994:Q2	N/A	N/A	−0.01	−0.04	−1.30	−0.85	−1.94	N/A
1994:Q3	N/A	N/A	0.47	0.07	−0.14	−1.52	−1.04	−2.00
1994:Q4	N/A	N/A	0.39	0.81	0.45	0.10	−1.28	−0.81
1995:Q1	N/A	N/A	−0.10	0.37	1.24	0.96	0.53	−0.81
1995:Q2	N/A	N/A	N/A	−0.60	−0.10	1.14	0.92	0.46
1995:Q3	N/A	N/A	−0.51	−0.79	−1.17	−0.64	0.88	0.70
1995:Q4	N/A	0.09	−0.39	−1.15	−1.39	−1.66	−1.11	0.60
1996:Q1	N/A	N/A	−0.63	−1.26	−2.07	−2.26	−2.44	−1.88
1996:Q2	N/A	0.19	−0.29	−1.21	−1.86	−2.65	−2.81	−2.96
1996:Q3	N/A	N/A	0.05	−0.60	−1.70	−2.26	−3.03	−3.19
1996:Q4	N/A	0.26	N/A	0.02	−0.64	−1.77	−2.23	−2.99
1997:Q1	N/A	N/A	0.20	−0.48	−0.18	−0.77	−1.88	−2.24
1997:Q2	N/A	N/A	−0.43	0.19	−0.54	−0.17	−0.70	−1.76
1997:Q3	N/A	0.65	N/A	−0.01	0.53	−0.14	0.27	−0.19
1997:Q4	N/A	0.28	0.64	0.72	0.11	0.53	−0.09	0.35
1998:Q1	N/A	−0.01	0.30	0.58	0.68	0.12	0.39	−0.18

(continued)

Table 11. (Continued)

Forecast Error	$r(t-1,t)$	$r(t-2,t)$	$r(t-3,t)$	$r(t-4,t)$	$r(t-5,t)$	$r(t-6,t)$	$r(t-7,t)$	$r(t-8,t)$
1998:Q2	N/A	0.33	0.11	0.34	0.57	0.67	0.14	0.29
1998:Q3	N/A	0.56	0.81	0.40	0.51	0.70	0.77	0.27
1998:Q4	N/A	−0.35	0.15	0.35	−0.14	−0.12	0.03	0.08
1999:Q1	N/A	−0.41	−1.27	−0.84	−0.70	−1.24	−1.29	−1.16
1999:Q2	N/A	0.40	−0.59	−1.49	−1.21	−1.10	−1.67	−1.78
1999:Q3	N/A	0.18	0.38	−0.44	−1.39	−1.24	−1.14	−1.75
1999:Q4	N/A	0.51	0.51	0.69	0.12	−0.86	−0.81	−0.73
2000:Q1	N/A	0.50	0.77	0.93	1.15	0.77	−0.22	−0.26
2000:Q2	N/A	−0.07	0.21	0.55	0.98	1.30	1.04	0.07
2000:Q3	N/A	−0.14	−0.29	0.00	0.25	0.92	1.34	1.14
2000:Q4	N/A	0.05	−0.36	−0.40	−0.09	0.07	0.89	1.39
2001:Q1	N/A	0.21	−0.22	−0.58	−0.57	−0.27	−0.16	0.73
2001:Q2	N/A	0.02	−0.16	−0.76	−1.07	−1.07	−0.77	−0.70
2001:Q3	N/A	0.15	−0.11	−0.42	−1.04	−1.31	−1.36	−1.05
2001:Q4	N/A	−0.43	−0.66	−0.91	−1.23	−1.80	−2.03	−2.15
2002:Q1	N/A	0.23	−0.83	−0.95	−1.14	−1.46	−1.96	−2.15

(continued)

Table 11. (Continued)

Forecast Error	$r(t-1,t)$	$r(t-2,t)$	$r(t-3,t)$	$r(t-4,t)$	$r(t-5,t)$	$r(t-6,t)$	$r(t-7,t)$	$r(t-8,t)$
2002:Q2	N/A	−0.01	0.08	−1.01	−1.02	−1.15	−1.44	−1.89
2002:Q3	N/A	−0.17	−0.41	−0.14	−1.12	−1.08	−1.14	−1.42
2002:Q4	N/A	0.26	−0.59	−0.68	−0.32	−1.19	−1.12	−1.13
2003:Q1	N/A	0.13	0.05	−1.05	−1.00	−0.62	−1.37	−1.29
2003:Q2	N/A	0.37	0.07	−0.23	−1.37	−1.22	−0.84	−1.49
2003:Q3	N/A	0.19	0.26	−0.23	−0.66	−1.74	−1.51	−1.16
2003:Q4	N/A	0.29	0.38	0.41	−0.24	−0.72	−1.72	−1.44
2004:Q1	N/A	−0.05	0.31	0.63	0.62	−0.11	−0.61	−1.53
2004:Q2	N/A	0.27	0.04	0.38	0.87	0.84	0.09	−0.42
2004:Q3	N/A	0.23	0.55	0.30	0.62	1.21	1.16	0.41
2004:Q4	N/A	−0.05	0.07	0.56	0.26	0.57	1.21	1.15

Note: Error is Actual minus Forecast.

References

Adolfson, M., M. K. Andersson, J. Lindé, M. Villani, and A. Vredin. 2007. "Modern Forecasting Models in Action: Improving Macroeconomic Analyses at Central Banks." *International Journal of Central Banking* 3 (4): 111–44.

Anderson, N., and J. Sleath. 1999. "New Estimates of the UK Real and Nominal Yield Curves." *Quarterly Bulletin* (Bank of England) (November): 384–92.

———. 2001. "New Estimates of the UK Real and Nominal Yield Curves." Working Paper No. 126, Bank of England.

Andersson, M., and B. Hofmann. 2009. "Gauging the Effectiveness of Quantitative Forward Guidance: Evidence from Three Inflation Targeters." ECB Working Paper No. 1098 (October).

Brooke, M., N. Cooper, and C. Scholtes. 2000. "Inferring Market Interest Rate Expectations from Money Market Rates." *Quarterly Bulletin* (Bank of England) (November): 392–402.

Ferrero, G., and A. Nobili. 2009. "Futures Contract Rates as Monetary Policy Forecasts." *International Journal of Central Banking* 5 (2): 109–45.

Ferrero, G., and A. Secchi. 2009. "The Announcement of Monetary Policy Intentions." Working Paper No. 720 (September), Bank of Italy.

Goodhart, C. A. E. 2009. "The Interest Rate Conditioning Assumption." *International Journal of Central Banking* 5 (2): 85–108.

Goodhart, C. A. E., and W. B. Lim 2008. "Interest Rate Forecasts: A Pathology." Discussion Paper No. 612 (June), Financial Markets Group, London School of Economics.

———. 2009. "The Value of Interest Rate Forecasts?" Special Paper No. 185 (April), Financial Markets Group, London School of Economics.

Joyce, M., J. Relleen, and S. Sorensen. 2007. "Measuring Monetary Policy Expectations from Financial Market Instruments." Prepared for ECB workshop on The Analysis of the Money Market: Role, Challenges and Implications from the Monetary Policy Perspective, Frankfurt, November 14–15.

Lange, J., B. Sack, and W. Whitesell. 2003. "Anticipation of Monetary Policy in Financial Markets." *Journal of Money, Credit, and Banking* 35 (6): 889–909.

Lildholdt, P., and A. Y. Wetherilt. 2004. "Anticipations of Monetary Policy in UK Financial Markets." Working Paper No. 241, Bank of England.

Mincer, J., and V. Zarnowitz. 1969. "The Evaluation of Economic Forecasts." In *Economic Forecasts and Expectation*, ed. J. Mincer, 3–46. New York: Columbia University Press.

Morris, S., and H. S. Shin. 2002. "Social Value of Public Information." *American Economic Review* 92 (5): 1521–34.

Rudebusch, G. D. 1995. "Federal Reserve Interest Rate Targeting, Rational Expectations, and the Term Structure." *Journal of Monetary Economics* 35 (2): 245–74.

———. 2002. "Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia." *Journal of Monetary Economics* 49 (6): 1161–87.

———. 2006. "Monetary Policy Inertia: Fact or Fiction?" *International Journal of Central Banking* 2 (4): 85–135.

———. 2007. "Monetary Policy Inertia and Recent Fed Actions." Federal Reserve Bank of San Francisco Economic Letter No. 2007-03 (January 26).

Svensson, L. E. O. 2006. "Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro-Transparency, Not Con." *American Economic Review* 96 (1): 448–52.

Endogenous Exposure to Systemic Liquidity Risk*

Jin Cao[a] and Gerhard Illing[b,c]
[a]Norges Bank
[b]Department of Economics, University of Munich
[c]CESifo

Traditionally, aggregate liquidity shocks are modeled as exogenous events. This paper analyzes the adequate policy response to endogenous exposure to systemic liquidity risk. We analyze the feedback between lender-of-last-resort policy and incentives of private banks, determining the aggregate amount of liquidity available. We show that imposing minimum liquidity standards for banks ex ante is a crucial requirement for sensible lender-of-last-resort policy. In addition, we analyze the impact of equity requirements and narrow banking, in the sense that banks are required to hold sufficient liquid funds so as to pay out in all contingencies. We show that both policies are strictly inferior to imposing minimum liquidity standards ex ante combined with lender-of-last-resort policy.

JEL Codes: E5, G21, G28.

"The events earlier this month leading up to the acquisition of Bear Stearns by JP Morgan Chase highlight the importance of liquidity management in meeting obligations during stressful market conditions. . . . The fate of Bear Stearns was the result of a lack of confidence, not a lack of capital. . . . At all

*The authors thank Antoine Martin, Rafael Repullo, Jean-Charles Rochet, Marti G. Subrahmanyam, Ted Temzelides, co-editor Douglas Gale, and seminar participants at various conferences for useful comments. Jin Cao gratefully acknowledges financial support at the Munich Graduate School of Economics (MGSE) from the German Research Foundation (DFG) through GRK 801. The views expressed should not be interpreted as reflecting those of Norges Bank.
Corresponding author: Gerhard Illing, Seminar für Markoökonomie, Ludwig-Maximilians-Universität München, Ludwigstrasse 28/012 (Rgb.), D-80539 Munich, Germany. Tel.: +49 89 2180 2126; fax: +49 89 2180 13521. Author e-mail addresses: jin.cao@gmail.com (Jin Cao), illing@lmu.de (Gerhard Illing).

times until its agreement to be acquired by JP Morgan Chase during the weekend, the firm had a capital cushion well above what is required to meet supervisory standards calculated using the Basel II standard."

(SEC Chairman Christopher Cox, "Letter to Basel Committee in Support of New Guidance on Liquidity Management," March 20, 2008)

"Bear Stearns never ran short of capital. It just could not meet its obligations. At least that is the view from Washington, where regulators never stepped in to force the investment bank to reduce its high leverage even after it became clear Bear was struggling last summer. Instead, the regulators issued repeated reassurances that all was well. Does it sound a little like a doctor emerging from a funeral to proclaim that he did an excellent job of treating the late patient?"

(Floyd Norris, "The Regulatory Failure Behind the Bear Stearns Debacle," *New York Times*, April 4, 2008)

1. Introduction

Before the financial crisis, financial markets seemed to have been awash with excessive liquidity. But suddenly, in August 2007, liquidity dried up nearly completely as a response to doubts about the quality of subprime mortgage-backed securities. Despite massive central bank interventions, the liquidity freeze did not melt away, but rather spread slowly to other markets such as those for auction-rate bonds. On March 16, 2008, the investment bank Bear Stearns, which (according to the Securities and Exchange Commission chairman) was adequately capitalized even a week before, had to be rescued via a Fed-led takeover by JP Morgan Chase.

Following the turmoil on financial markets, there has been a strong debate about the adequate policy response. Some have warned that central bank actions may encourage dangerous moral hazard behavior of market participants in the future. Others instead criticized central banks for responding far too cautiously. The most prominent voice has been Willem Buiter, who—jointly with Ann Sibert—right from the beginning of the crisis in August 2007 strongly pushed the idea that in times of crises, central banks should act as market maker of last resort (see Buiter and Sibert 2007). As an adaptation of the Bagehot principle to modern times with globally integrated financial systems, central banks should actively purchase and sell illiquid private-sector securities and so play a key role in assessing and pricing credit risk. In his *Financial Times* blog "Maverecon," Willem Buiter stated the intellectual arguments behind such a policy very clearly on December 13, 2007:

> Liquidity is a public good. It can be managed privately (by hoarding inherently liquid assets), but it would be socially inefficient for private banks and other financial institutions to hold liquid assets on their balance sheets in amounts sufficient to tide them over when markets become disorderly. They are meant to intermediate short maturity liabilities into long maturity assets and (normally) liquid liabilities into illiquid assets. Since central banks can create unquestioned liquidity at the drop of a hat, in any amount and at zero cost, they should be the liquidity providers of last resort, both as lender of last resort and as market maker of last resort. There is no moral hazards as long as central banks provide the liquidity against properly priced collateral, which is in addition subject to the usual "liquidity haircuts" on this fair valuation. The private provision of the public good of emergency liquidity is wasteful. It's as simple as that.

Buiter's statement represents the prevailing mainstream view that there is no moral hazard risk as long as the Bagehot principle is followed as best practice in liquidity management.

According to Goodfriend and King (1988), a lender-of-last-resort policy should target liquidity provision to the market, but not to specific banks. Central banks should "lend freely at a high rate against good collateral." This way, public liquidity support is supposed to be targeted toward solvent yet illiquid institutions, since insolvent financial institutions should be unable to provide adequate collateral to secure lending. This paper wants to challenge the view that a policy following the Bagehot principle does not create moral hazard. The key argument is that this view neglects the endogeneity of aggregate liquidity risk. Starting with Allen and Gale (1998) and Holmström and Tirole (1998), there have been quite a few models recently analyzing private and public provision of liquidity. But in most of these models, exposure to aggregate systemic risk is assumed to be exogenous.

In Holmström and Tirole (1998), for instance, liquidity shortages arise when financial institutions and industrial companies scramble for and cannot find the cash required to meet their most urgent needs or undertake their most valuable projects. They show that credit lines from financial intermediaries are sufficient for implementing the socially optimal (second-best) allocation, as long as there is no aggregate uncertainty. In the case of aggregate uncertainty, however, the private sector cannot satisfy its own liquidity needs, so the existence of liquidity shortages vindicates the injection of liquidity by the government. In their model, the government can provide (outside) liquidity by committing future tax income to back up the reimbursements.

In the model of Holmström and Tirole (1998), the lender of last resort indeed provides a free lunch: Public provision of liquidity in the presence of aggregate shocks is a pure public good, with no moral hazard involved. The reason is that the probability for being hit by an aggregate shock is not affected by the amount of investment in liquid assets carried out by the private sector. The same holds in Allen and Gale (1998), even though they analyze a quite different mechanism for public provision of liquidity: the adjustment of the price level in an economy with nominal contracts. We adopt Allen and Gale's mechanism, but we endogenize the exposure of financial intermediaries to aggregate (systemic) liquidity risk.

The basic idea of our model is fairly straightforward: Financial intermediaries choose the share invested in projects which might turn out to be illiquid. We model (real) illiquidity in the following way: Some fraction of those projects turns out to be realized late. The aggregate share of late projects is endogenous, since it depends on the incentives of financial intermediaries to invest in those illiquid projects. When intermediaries would invest only in liquid assets, they would never be hit by shocks affecting illiquid projects. The larger the share invested in those assets, however, the higher the exposure to aggregate liquidity risk. This endogeneity allows us to capture the feedback from liquidity provision to risk-taking incentives of financial intermediaries. We show that the share invested in illiquid projects rises endogenously with central bank liquidity provision: The anticipation of unconditional central bank liquidity provision encourages excessive risk taking (moral hazard). It turns out that in the absence of liquidity requirements, there will be overinvestment in risky activities, creating excessive exposure to systemic risk.

In contrast to what the Bagehot principle suggests, unconditional provision of liquidity to the market (lending of central banks against good collateral) is exactly the wrong policy: It distorts incentives of banks to provide sufficient private liquidity, thus reducing investors' payoff. In our model, we concentrate on pure illiquidity risk: There will never be insolvency unless triggered by illiquidity (by a bank run). Illiquid projects promise a higher, yet possibly delayed, return. Relying on sufficient liquidity provided by the market (or by the central bank), financial intermediaries are inclined to invest more heavily in high-yielding, but illiquid, long-term projects. A central bank's liquidity provision, helping to prevent bank runs with inefficient early liquidation, encourages banks to invest more in illiquid assets. At first sight, this seems to work fine, even if systemic risk increases: After all, public insurance against aggregate risks should allow agents to undertake more profitable activities with higher social return. As long as public insurance is a free lunch, there is nothing wrong with providing such a public good.

The problem, however, is that due to limited liability some banks will be encouraged to free-ride on liquidity provision. This competition will force the other banks to reduce their efforts for liquidity provision, too. Chuck Prince, at that time chief executive of Citigroup, stated the dilemma posed in fairly poetic terms on July 10, 2007 in an infamous interview with *Financial Times*[1]:

> When the music stops, in terms of liquidity, things will be complicated. But as long as the music is playing, you've got to get up and dance. We're still dancing.

[1]The key problem is best captured by the following remark about Citigroup in the *New York Times* report "Treasury Dept. Plan Would Give Fed Wide New Power" on March 29, 2008: "Mr. Frank said he realized the need for tighter regulation of Wall Street firms after a meeting with Charles O. Prince III, then chairman of Citigroup. When Mr. Frank asked why Citigroup had kept billions of dollars in 'structured investment vehicles' off the firm's balance sheet, he recalled, Mr. Prince responded that Citigroup, as a bank holding company, would have been at a disadvantage because investment firms can operate with higher debt and lower capital reserves."

The dancing banks simply enjoy liquidity provided in good states of the world and just disappear (go bankrupt) in bad states. The incentive of financial intermediaries to free-ride on liquidity in good states results in excessively low liquidity in bad states. Even worse, as long as they do not suffer runs, "dancing" banks can always offer more attractive collateral in bad states—so they are able to outbid prudent banks in a liquidity crisis. For that reason, the Bagehot principle, rather than providing correct incentives, is the wrong medicine in modern times with a shadow banking system relying on liquidity being provided by other institutions.

This paper extends a model developed in Cao and Illing (2008). In that paper we did not allow for banks holding equity, so we could not analyze the impact of equity requirements. As we will show, imposing equity requirements can be inferior even to the outcome of a mixed-strategy equilibrium with free-riding (dancing) banks. In contrast, imposing binding liquidity requirements ex ante combined with lender-of-last-resort policy ex post is able to implement the second-best outcome. In our model, it yields a strictly superior outcome compared with imposing equity requirements. We also prove that "narrow banking" (banks being required to hold sufficient equity so as to be able to pay out demand deposits in all states of the world) is inferior to ex ante liquidity regulation.

Allen and Gale (2007, p. 213f) notice that the nature of market failure leading to systemic liquidity risk is not yet well understood. They argue that "a careful analysis of the costs and benefits of crises is necessary to understand when intervention is necessary." In this paper, we try to fill this gap, providing a cost/benefit analysis of different forms of banking regulation to better understand what type of intervention is required. We explicitly compare the impact both of liquidity and of equity requirements. To the best of our knowledge, this is the first paper providing such an analysis.

The paper is organized as follows. In section 2, we present the structure of the model with real deposit contracts and characterize the central planner's constrained efficient solution and the market equilibrium. In section 3 we introduce a central bank and show that in an economy with nominal deposit contracts, lender-of-last-resort policy eliminates bank runs but is subject to the time-inconsistency problem. We show that ex ante liquidity regulation, combined with lender-of-last-resort policy, can implement the constrained efficient solution. The effectiveness of imposing equity requirements is analyzed in section 4. Section 5 concludes.

2. The Structure of the Model

2.1 The Agents, Time Preferences, and Technology

In this economy, there are three types of agents: investors, banks (run by bank managers), and entrepreneurs. All agents are risk neutral. The economy extends over three periods, $t = 0, 1, 2$, and the details of timing will be explained later. We assume the following:

(i) There is a continuum of investors, each initially (at $t = 0$) endowed with one unit of resources. The resource can be either stored (with a gross return equal to 1) or invested in the form of bank deposits.

(ii) There is a finite number N of active banks engaged in Bertrand competition, competing for investors' deposits. Using these deposits, the banks as financial intermediaries can fund projects of entrepreneurs.

(iii) There is a continuum of entrepreneurs. There are two types of them (denoted by i, $i = 1, 2$), characterized by their project returns R_i:
- Projects of type 1 (safe projects) are realized early at period $t = 1$ with a safe return $R_1 > 1$.
- Projects of type 2 (risky projects) give a higher return $R_2 > R_1 > 1$. With probability p, these projects will also be realized at $t = 1$, but they may be delayed (with probability $1 - p$) until $t = 2$. Therefore, in the aggregate, the share p of type 2 projects will be realized early. The aggregate share p, however, is not known at $t = 0$. It will only be revealed between periods 0 and 1 at some intermediate period; call it $t = \frac{1}{2}$. In the following, we are interested in the case of aggregate shocks. We model them in the simplest way: The aggregate share of type 2 projects realized early, p, can take on just two values—either p_H or p_L with $p_H > p_L$. The "good" state with a high share of early type 2 projects p_H, i.e., the state with plenty of liquidity, will

be realized with probability π. In the following, we assume that $1 < p_s R_2 < R_1$ $(s \in \{H, L\})$ to focus on the relevant case (to be explained later).

Investors are impatient: They want to consume early (at $t = 1$). In contrast, both entrepreneurs and bank managers are indifferent between consuming early ($t = 1$) or late ($t = 2$).

Focusing on the case of liquidity constraints being binding, we assume that resources of investors are scarce in the sense that there are more projects of each type available than the aggregate endowment of investors. Thus, in a first-best market economy (in the absence of commitment problems, as explained in the next paragraph), total surplus would go to the investors. They would simply put all their funds in early projects and capture the full return. We take this frictionless market outcome as our reference point and seek to minimize the distance in terms of the investors' welfare between this reference point and the equilibrium outcome under various policies. Hold-up problems prevent realization of the frictionless market outcome, creating a demand for liquidity. Since there is a market demand for liquidity only if investors' funds are the limiting factor, we choose the investors' payoff as the policymaker's objective and concentrate on deviations from this market outcome. With investors' payoff as the relevant criterion, we analyze those equilibria coming closest to implementing the frictionless market outcome.

Due to hold-up problems as modeled in Hart and Moore (1994), or Holmström and Tirole (1997), entrepreneurs can only commit to pay a fraction $\gamma < 1$ of their return with $\gamma R_i > 1$. Banks as financial intermediaries can pool investment; they have superior collection skills (a higher γ, which justifies their role as intermediaries). In the following, we also assume that $p_s \leq \gamma$ $(s \in \{H, L\})$ to concentrate on the relevant case that investors care about investment in liquid projects (see section 2.4). Following Diamond and Rajan (2001), banks offer deposit contracts with a fixed payment d_0 payable at any time after $t = 0$ as a credible commitment device not to abuse their collection skills. The threat of a bank run disciplines bank managers to fully pay out all available resources pledged in the form of bank deposits. Deposit contracts, however, introduce a fragile structure into the economy: Whenever investors have doubts about their bank's liquidity (the ability to pay investors the promised amount

d_0 at $t = 1$), they run on the bank at the intermediate date, forcing the bank to liquidate all its projects (even those funding entrepreneurs with safe projects) at high costs: Early liquidation of projects gives only the inferior return $c < 1$. In the following, we do not consider pure sunspot bank runs of the Diamond and Dybvig type. Instead, we concentrate on the runs happening if liquid funds are not sufficient to pay out investors.

2.2 Timing and Events

At date $t = 0$, banks competing for funds offer deposit contracts with payment d_0 which maximize the expected return of investors. Banks compete by choosing the share α of deposits invested in type 1 projects, taking their competitors' choice as given. Investors have rational expectations about each bank's default probability; they are able to monitor all banks' investment. Remember that, at this stage, the share p of type 2 projects that will be realized early is not yet known.

At date $t = \frac{1}{2}$, the value of p is revealed, as is the expected return of the banks at $t = 1$. A bank will experience a run if it cannot meet the investors' demand. If this happens, all the assets—even the safe projects—have to be liquidated.

Those banks which do not suffer a run trade with early entrepreneurs in a perfectly competitive market for liquidity at $t = 1$, clearing at interest rate r. Note that because of the hold-up problem, entrepreneurs retain a rent—their share $(1 - \gamma)R_i$. Since early entrepreneurs are indifferent between consuming at $t = 1$ or $t = 2$, they are willing to provide liquidity (using their rent to deposit at banks at $t = 1$ at the market rate r). Banks use the liquidity provided to pay out investors. In this way, impatient investors can profit indirectly from the investment in high-yielding long-term projects. So banking allows the transformation between liquid claims and illiquid projects.

At date $t = 2$, the banks collect the return from the late projects and pay back the early entrepreneurs at the predetermined interest rate r.

Note that the aggregate liquidity available at date $t = 1$ depends on the total share of funds, α, invested in liquid type 1 projects at

Figure 1. Timing and Payoff Structure, When Banks Are Liquid

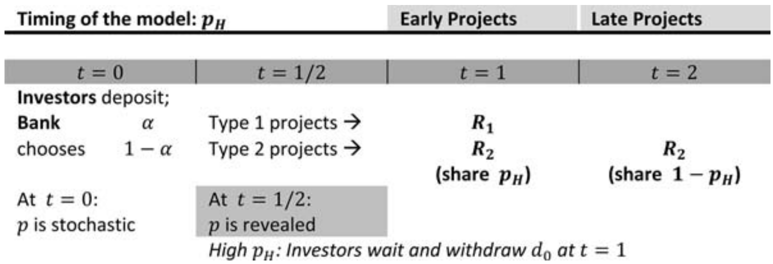

Figure 2. Timing and Payoff Structure, When Banks Are Illiquid

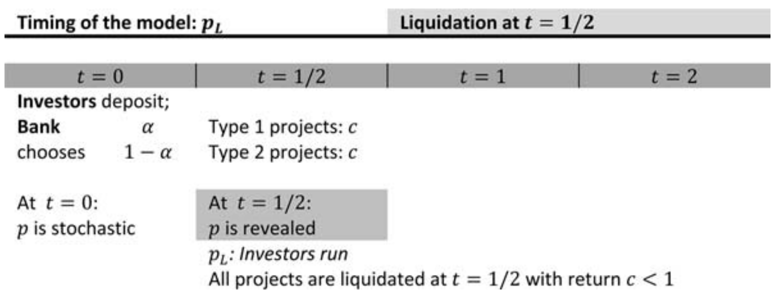

date $t = 0$. As long as the banks are liquid, the payoff structure is described as in figure 1. But if α is so low that the banks cannot honor deposits when p_L occurs, investors will run at $t = \frac{1}{2}$. The payoff in that case is captured in figure 2.

2.3 The Central Planner's Constrained Efficient Solution

We first analyze the problem of a central planner maximizing the investors' payoff. This provides the reference point for the market equilibrium with banks as financial intermediaries characterized in the next section. Investors being impatient, the central planner would choose the share invested in illiquid projects so as to maximize the resources available to investors at period 1. Since $p_s R_2 < R_1$, in the absence of hold-up problems, the planner would invest only in liquid type 1 projects, this way maximizing resources available at period 1. But due to the hold-up problem caused by entrepreneurs, the central planner can implement only a constrained efficient solution. If the central planner had unlimited taxation authority, he or she could eliminate the hold-up problem completely by taxing the entrepreneurs' rent and redistributing the resources to the investors. Again, all resources would be invested only in liquid type 1 projects, and the entrepreneurs' rents would be transferred to the investors in period 1.

Obviously, allowing for non-distortionary taxation biases the comparison between the market and the planner's solution, giving the planner an unfair advantage. Effectively, redistribution via lump-sum taxation would make both hold-up and liquidity constraints non-binding, assuming the relevant issues away. To make the planner's constrained optimization problem interesting, we assume that non-distortionary taxation is not feasible in period 1. In order to impose sensible restrictions, we take private endowments as a binding constraint and assume that the entrepreneur has to receive an equivalent compensation when he or she is asked to give up resources in period 1. Being indifferent between consuming at $t = 1$ and $t = 2$, the entrepreneur needs to be compensated by an appropriate transfer in period 2. In order not to distort the comparison in favor of banks, we furthermore assume that the planner has the same collection skills (the same γ) as financial intermediaries.

Given these constraints, the constrained efficient solution is characterized in the following proposition.

PROPOSITION 1. *The optimal solution for the central planner's problem is as follows:*

(i) If there is no aggregate risk, i.e., when p_s is known at $t = 0$, the planner invests the share $\alpha_s = \frac{\gamma - p_s}{\gamma - p_s + (1-\gamma)\frac{R_1}{R_2}}$ $(s \in \{H, L\})$ in liquid projects and the investors' return is maximized at $\gamma E[R_s] = \gamma[\alpha_s R_1 + (1 - \alpha_s)R_2]$.

(ii) In the presence of aggregate risk, the central planner implements the following state-contingent strategy, depending on

the probability π for p_H being realized: The planner invests the share α_H in liquid projects as long as $\overline{\pi}'_2 = \frac{\gamma E[R_L]-\kappa}{\gamma E[R_H]-\kappa} \leq \pi \leq 1$ with $\kappa = \alpha_H R_1 + (1-\alpha_H)p_L R_2$, and the share $\alpha_L > \alpha_H$ for $0 \leq \pi < \overline{\pi}'_2$.

Proof. See appendix 1.

The first part of the proposition says that if p is known, the planner simply chooses α so as to maximize the investors' return. The second part says that if p is unknown, the planner faces a trade-off: the investors' return is maximized under p_H if the planner chooses α_H, but will be low if p_L is realized; the investors' return is maximized under p_L if the planner chooses α_L, but will be low if p_H is realized. So the optimal solution depends on the likelihood of p_H, that is, on π. When π is high enough, the planner will choose α_H; otherwise, he or she will pick α_L.

Obviously, hold-up and liquidity constraints are bound to have a distributional impact: If resources were taken away from investors in the initial period and redirected toward type 2 entrepreneurs, the commitment problem would no longer be relevant, nor would the need for liquidity provision. Even though such a reallocation would result in higher aggregate resources (all funds being invested in high-return projects), it would yield inferior payoff to investors. Since $p_s R_2 < R_1$, investing less than α_s in liquid projects reduces resources available in period 1 and so makes investors worse off.

In contrast to our modeling strategy, Holmström and Tirole (1998) assume that the lender of last resort has unlimited power to tax real resources and so is always able to redistribute resources ex post. This assumption, however, effectively makes liquidity constraints non-binding: The central planner can always redistribute resources ex post in such a way as to make them irrelevant. The planner could simply redirect resources to the constrained agents (and potentially compensate the unconstrained). Interestingly, in our model, giving the planner taxation power in period 2 cannot help to improve upon the investors' allocation: The investors being impatient, any redistribution from illiquid projects realized late at $t = 2$ is simply not feasible.

2.4 The Market Equilibrium

Let us now characterize the market equilibrium with banks as financial intermediaries. First, let us again start with the simplest case with no aggregate uncertainty; i.e., the share p of type 2 projects realized early is known at $t = 0$. The market equilibrium of the model is characterized by bank i's strategic profile (α_i, d_{0i}), $\forall i \in \{1, \ldots, N\}$ such that

- Bank i's profit is maximized by

$$\alpha_i = \arg \max_{\alpha_i \in [0,1]} \gamma \left\{ \alpha_i R_1 + (1 - \alpha_i) \left[pR_2 + \frac{(1-p)R_2}{r} \right] \right\}. \tag{1}$$

  Bank i chooses the share of liquid projects α_i so as to maximize expected discounted returns;
- Bank i makes zero profit from offering deposit contract d_{0i},

$$d_{0i} = \max_{\alpha_i \in [0,1]} \gamma \left\{ \alpha_i R_1 + (1 - \alpha_i) \left[pR_2 + \frac{(1-p)R_2}{r} \right] \right\}. \tag{2}$$

  Investors deposit their funds at those banks offering the highest return. Thus, with Bertrand competition in the deposit market, the deposit rate d_{0i} offered to investors in equilibrium will be equal to expected returns, maximizing resources available at period 1;
- The market interest rate is determined in the following way:

  (i) In equilibrium, all resources available at $t = 1$ will be paid out to investors, so $d_{0i} = \alpha_i R_1 + (1 - \alpha_i)pR_2$. Banks receive funds $\gamma[\alpha_i R_1 + (1 - \alpha_i)pR_2]$ from those projects realized early. In addition, early entrepreneurs are willing to provide liquidity at $t = 1$ (depositing their rent at the market rate $r \geq 1$) to solvent banks that are able to meet their liabilities to the investors, that is, to banks with $d_{0i} \leq \gamma \left\{ \alpha_i R_1 + (1 - \alpha_i) \left[pR_2 + \frac{(1-p)R_2}{r} \right] \right\}$. So the liquidity supplied by early entrepreneurs is $(1 - \gamma)[\alpha_i R_1 + (1 - \alpha_i)pR_2]$ as long as bank i is expected to stay solvent—that is, as long as it is able to pay out early entrepreneurs at the market rate r at $t = 2$ from its late project's return $\gamma(1-\alpha_i)(1-p)R_2$.

Furthermore, as the market clearing condition, aggregate liquidity supply and demand at $t = 1$ have to be equal, given that banks stay solvent at the interest rate $r \geq 1$: $\sum_{i=1}^{N} r(1-\gamma)[\alpha_i R_1 + (1-\alpha_i)pR_2] = \sum_{i=1}^{N} \gamma(1-\alpha_i)(1-p)R_2$;

(ii) Finally, when there is excess liquidity supply at $t = 1$, i.e., when total intermediate output exceeds the payoff promised to the investors, $r = 1$.

If there is no aggregate uncertainty, the market equilibrium with $r = 1$ is equivalent to the solution of the social planner's problem: Banks will invest such that—on aggregate—they are able to fulfill investors' claims in period 1, so there will be no run.

PROPOSITION 2. *If there is no aggregate uncertainty, the allocation in the market equilibrium with* $r = 1$ *is identical to the solution of the social planner's problem, characterized by the following:*

- *All banks set* $\alpha = \frac{\gamma-p}{\gamma-p+(1-\gamma)\frac{R_1}{R_2}}$.
- *The market interest rate* $r = 1$.

Proof. See appendix 1.

The proposition says that in the absence of aggregate uncertainty, the banks will choose α (the share invested in liquid projects) so as to maximize the depositors' return and to stay solvent at $t = 1$, given that entrepreneurs are willing to provide liquidity at that time. This coincides with the solution of the social planner's problem. Since $R_1 > pR_2$ and $\gamma > p$, α will be strictly positive in equilibrium. For given p, there is a unique α maximizing resources available for investors at $t = 1$. A bank investing less than this value of α would not be able to pay out the amounts promised to investors at $t = 1$ and thus would experience a run at $t = \frac{1}{2}$. A bank investing more than α would be outbid by competitors offering a higher d_{0i}. Note that α is decreasing in p: The larger the share p of type 2 projects realized early, the less need for investment in liquid type 1 projects. For $p > \gamma$, liquid projects are dominated by the risky ones, so there would be no demand for liquid projects at $t = 0$. Similarly, there would be no demand for liquid projects at $t = 0$ either when $R_1 < pR_2$. Since liquidity is not an issue for these cases, they are ruled out by assumption.

It becomes tricky to find the market equilibrium when there is aggregate uncertainty. Let us briefly sketch the market equilibrium in the following proposition.

PROPOSITION 3. *When there is aggregate uncertainty,*

(i) there is a symmetric pure-strategy equilibrium such that all banks set $\alpha = \alpha_H$ *for all* $\overline{\pi}_2 < \pi \leq 1$*, with* $\overline{\pi}_2 = \frac{\gamma E[R_L]-c}{\gamma E[R_H]-c}$ *and* $E[R_s] = \alpha_s R_1 + (1 - \alpha_s)R_2$ $(s \in \{H, L\})$*;*

(ii) there is a symmetric pure-strategy equilibrium such that all banks set $\alpha = \alpha_L$ *for all* $0 \leq \pi < \overline{\pi}_1$*, with* $\overline{\pi}_1 = \frac{\gamma E[R_L]-c}{\gamma R_2 - c}$*; and*

(iii) there exists no symmetric pure-strategy equilibrium for all $\overline{\pi}_1 \leq \pi \leq \overline{\pi}_2$*. However, there exists a unique equilibrium in mixed strategies such that*

(a) at $t = 0$*, with probability* θ *a bank chooses to be a free-riding bank that sets* $\alpha = 0$ *and with probability* $1 - \theta$ *a bank chooses to be a prudent bank that sets* $0 < \alpha_s^* < \alpha_L$*; and*

(b) in the mixed-strategy equilibrium, investors are worse off than if all banks would coordinate on the prudent (non-equilibrium) strategy α_L*.*

Proof. See appendix 1.

The intuition behind proposition 3 is as follows: With uncertainty about p, a bank seems to have just two options available—it may either invest so much in safe type 1 projects (α_L) that it will be able to pay out the investors all the time (that is, even if the bad state occurs), or it may invest just enough, α_H, so as to pay out investors only in the good state and experience a run in the bad state. If π is very high (close to 1), a bank should choose α_H—to reap the high yields in the good state, since the cost of the bank run in the bad state is rather low. Alternatively, if π is very low (close to 0), it always pays to be prepared for the worst case, so the bank should choose $\alpha_L > \alpha_H$ in safe projects. Since α_s $(s \in \{H, L\})$ is the share invested in safe projects with return R_1, the total payoff by choosing α_s is $E[R_s] = \alpha_s R_1 + (1-\alpha_s)R_2$, with $E[R_H] > E[R_L]$.

With a high share α_L of safe projects, the banks will be able to pay out investors in all states. There will never be a bank run. So, independent of π, the expected payoff for investors is $\gamma E[R_L]$. In contrast, with strategy α_H there will be a bank run in the bad state, giving just the bankruptcy payoff c with probability $1-\pi$. So the return to strategy α_H is $\pi\gamma E[R_H] + (1-\pi)c$, which is increasing in π. Investors get a higher payoff under α_H, if $\pi\gamma E[R_H] + (1-\pi)c > \gamma E[R_L]$ or

$$\pi > \overline{\pi}_2 = \frac{\gamma E[R_L] - c}{\gamma E[R_H] - c}.$$

For $\pi < \overline{\pi}_2$, the investors' payoff is higher with strategy α_L. But if all banks would choose strategy α_L, there will be excess liquidity at $t = 1$ if the good state occurs (with a large share of type 2 projects realized early). A bank anticipating this event has a strong incentive to invest all funds in type 2 projects, reaping the benefit of excess liquidity in the good state. As long as the music is playing, such a deviating bank gets up and dances. In the good state, such a free-riding bank can credibly rely on entrepreneurs' excess liquidity at $t = 1$, promising to pay back at $t = 2$ out of highly profitable projects. After all, at that stage, this bank, free-riding on liquidity, can offer a capital cushion with expected returns well above what prudent banks are able to promise. Of course, if the bad state happens, there is no excess liquidity. Liquidity dries up. The free-riding banks would just bid up the interest rates, urgently trying to get funds. Rational investors, anticipating that these banks will not succeed, will have already triggered a bank run on these banks at $t = \frac{1}{2}$.

As long as the free-riding banks are not supported in the bad state, they are driven out of the market, providing just the return c. Nevertheless, these banks can offer the return $\pi\gamma R_2 + (1-\pi)c$ as expected payoff for investors. Thus, a free-riding bank will be able to offer a higher expected return than a prudent bank, provided the probability π for the good state is not too low. The condition is

$$\pi > \overline{\pi}_1 = \frac{\gamma E[R_L] - c}{\gamma R_2 - c}.$$

Figure 3. Investors' Expected Return under Free-Riding Compared with the Case of No Free-Riding and the Central Planner's Solution

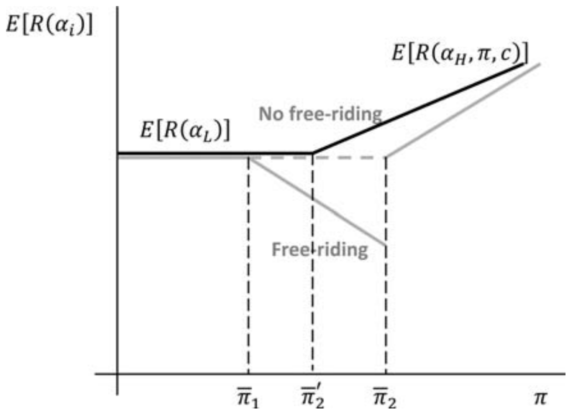

Since $R_2 > E[R_H]$, it pays to free-ride within the range $\overline{\pi}_1 \leq \pi < \overline{\pi}_2$.

Obviously, there cannot be an equilibrium in pure strategies within that range. As long as the music is playing, all banks would like to "get up and dance." But then, there would be no prudent bank left providing the liquidity needed to be able to free-ride. In the resulting mixed-strategy equilibrium, a proportion of banks behave prudently, investing some amount $\alpha_s^* < \alpha_L$ in liquid assets, whereas the rest free-ride on liquidity in the good state, choosing $\alpha = 0$. Prudent banks reduce α in order to cut down the opportunity cost of investing in safe projects. Interest rates and α_s^* adjust so that investors are indifferent between the two types of banks. At $t = 0$, both prudent and free-riding banks offer the same expected return to investors. The proportion of free-riding banks is determined by aggregate market clearing conditions in both states. Free-riding banks experience a run for sure in the bad state, but the high return in the good state R_2 compensates investors for that risk.

As shown in proposition 3, free-riding drives down the return for investors. They are definitely worse off than they would be if all banks coordinated on the prudent strategy α_L—similar to the inefficient mixed-strategy equilibrium in Allen and Gale (2004). The solid gray lines in figure 3 illustrate the investors' expected return in the market equilibrium; as a result of free-riding behavior, the effective return on deposits for investors deteriorates in the range

$\overline{\pi}_1 \leq \pi < \overline{\pi}_2$, compared with the outcome if all banks would coordinate (off equilibrium) on α_L as the dashed gray line shows.

Compared with the central planner's solution (the solid black line in figure 3), the investor's payoff is lower in the market equilibrium with banks as financial intermediaries for two different reasons: First, free-riding banks reduce the investor's payoff in the mixed-strategy equilibrium in the intermediate case. Second, for high values of π ($\overline{\pi}_2 < \pi \leq 1$), a representative bank, choosing α_H, accepts the risk of a bank run if the bad state occurs (with a low share p_L of illiquid projects realized early). If that state occurs, a bank run is triggered with inefficient liquidation, resulting in an inferior payoff $c < 1$. In the following sections, we will carefully analyze how different mechanism designs may help to raise the investor's payoff, bringing the market outcome closer to the constrained efficient solution as stated in section 2.3. In the next section, we show that in an economy with nominal deposit contracts, lender-of-last-resort policy is able to tackle the problem of bank runs, but at the same time aggravates banks' incentives for free riding.

3. Lender-of-Last-Resort Policy

3.1 Nominal Contracts and the Lender of Last Resort

A lender of last resort, usually the central bank, cannot create real liquidity in period 1. But a central bank can add nominal liquidity at the stroke of a pen. Following Allen and Gale (1998) and Diamond and Rajan (2006), assume from now on that deposit contracts are arranged in nominal terms. The bail-out mechanism of the central bank is similar to that in Allen, Carletti, and Gale (2010). Here is the timing of the model:

(i) At $t = 0$, the banks provide nominal deposit contracts to investors, promising a fixed nominal payment d_0 at $t = 1$. The central bank announces a minimum level $\underline{\alpha}$ of investment on safe projects, and only those banks that meet the requirement will be eligible for liquidity support in the time of a crisis.

(ii) At $t = \frac{1}{2}$, the banks decide whether to borrow liquidity from the central bank. If yes, the central bank will provide liquidity for the banks, provided they fulfill the requirement $\underline{\alpha}$.

(iii) At $t = 1$, the liquidity injection with the banks' illiquid assets as collateral is carried out so that the banks are able to honor their nominal contracts, which reduces the real value of deposits just to the amount of real resources available at that date.
(iv) At $t = 2$, the banks repay the central bank using the returns from the late projects, with gross nominal interest rate r^M agreed at $t = 1$.

As a new element in this extended model, we allow the central bank to impose minimal liquidity holdings in addition to lender-of-last-resort policy as a way to implement the allocation maximizing the investors' payoff.

3.2 Liquidity Regulation and Lender-of-Last-Resort Policy

With nominal contracts, the central bank's optimal policy as lender of last resort can be summarized in the following proposition.

PROPOSITION 4. *With nominal contracts, the central bank can act as lender of last resort. The central bank's optimal policy that maximizes the investors' return is as follows:*

(i) Set $\underline{\alpha} = \alpha_H$ *for all* $\pi \in [\overline{\pi}'_2, 1]$*, where* $\overline{\pi}'_2 = \frac{\gamma E[R_L]-\kappa}{\gamma E[R_H]-\kappa} < \overline{\pi}_2$ *and* $\kappa = \alpha_H R_1 + (1-\alpha_H)p_L R_2$*.*
(ii) Set $\underline{\alpha} = \alpha_L$ *for all* $\pi \in [0, \overline{\pi}'_2)$*.*
(iii) Set $r^M = 1$*.*

What's more, under such a policy bank runs are eliminated for the eligible banks; i.e., the eligible banks will not experience runs when p_L *is revealed.*

Proof. See appendix 1.

The investors' return is maximized when the banks get liquidity injection at the lowest cost, the central bank setting $r^M = 1$. With liquidity injection, bank runs are prevented when the bad state (with low payoffs at $t = 1$) occurs. Essentially, nominal deposits allow the central bank to implement state-contingent payoffs, and such a policy replicates the optimal allocation in the central planner's problem. This argument seems to confirm the view that the lender of last resort can indeed provide a free lunch, delivering a public good at no cost. It turns out, however, that the anticipation of these actions has an adverse impact on the amount of aggregate liquidity provided by the private sector, affecting endogenously the exposure to systemic risk.

PROPOSITION 5. *Assume that a market equilibrium exists, i.e.,* $\pi p_H R_2 + (1 - \pi) p_L R_2 \geq 1$. *If the central bank is willing to provide liquidity to the entire market in times of crisis, all banks have an incentive to free-ride, choosing* $\alpha = 0$, *and investors are made worse off.*

Proof. See appendix 1.

The reason for this surprising result is the following: If the central bank targets liquidity provision to the market instead of to specific banks, the optimal policy as stated by proposition 4 is not enforceable. Since we concentrate on the case of pure illiquidity risk, in our model all projects will certainly be realized at $t = 2$. So there is no doubt about solvency of the projects, unless insolvency is triggered by illiquidity. If the central bank follows the Bagehot principle and creates artificial liquidity at the drop of a hat—against allegedly good collateral—all private incentives to care about ex ante liquidity provision will be destroyed, exacerbating the moral hazard problem: The free-riding banks, investing all their funds in the projects with higher returns, can always get liquidity support and thus are able to offer more attractive terms to investors at $t = 0$. This drives out all the prudent banks and leaves the investors worse off.[2]

So what policy options should be taken? One might argue that a central bank should provide liquidity support only to prudent banks (conditional on banks having invested sufficiently in liquid assets). But such a commitment is simply not credible: As emphasized by Rochet (2004) and Cao and Illing (2008), there is a serious problem of dynamic consistency.

[2]In reality, there is no clear-cut distinction between insolvency and illiquidity. We leave it to future research to allow for the risk of insolvency. But we doubt that our basic argument will be affected.

Figure 4. Investors' Real Expected Return $E[R(\alpha_H, \pi, \kappa)]$ for the Case of Ex Ante Liquidity Regulation Combined with Ex Post Lender-of-Last-Resort Policy for High π, in Comparison with the Market Equilibrium, $E[R(\alpha_H, \pi, c)]$

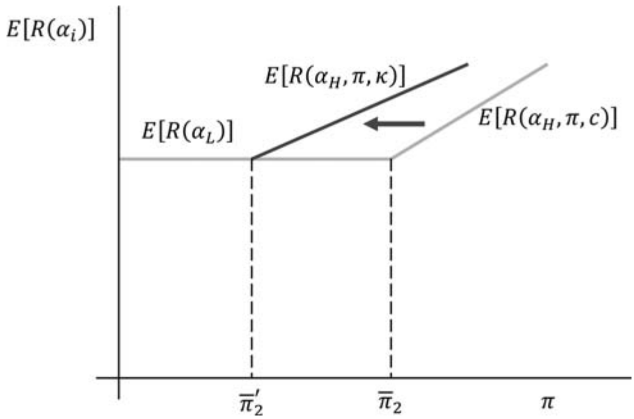

Rather than relying on an implausible commitment mechanism, the obvious solution is a mix of two instruments: ex ante liquidity regulation combined with ex post lender-of-last-resort policy. The second-best outcome from the investors' point of view needs to be implemented by the following policy: In a first step, a banking regulator has to impose ex ante liquidity requirements. Requesting minimum investment in liquid type 1 assets of at least α_L for $\pi < \overline{\pi}'_2$ and α_H for $\pi \geq \overline{\pi}'_2$ would give investors the highest expected payoff, as characterized in figure 4. When banks are not allowed to operate with insufficiently low liquidity holdings, there are no incentives for free-riding. For high values $\pi \geq \overline{\pi}'_2$, the central bank acts as lender of last resort in the bad state, eliminating costly bank runs. This raises the expected payoff for investors, even though it increases the range of parameter values with systemic risk.

In a quite different setting, using a framework with asymmetric information, Farhi and Tirole (2009) derive related results. They show that monetary policy (with the real interest rate as policy variable) faces a commitment problem. They also derive a role for a minimum liquidity ratio. Our setup shows that the key challenge for regulators and the central bank is to cope with incentives for financial intermediaries to free-ride on liquidity provision. Furthermore, it allows us to compare liquidity regulation with alternative mechanism designs. One might expect that imposing equity requirements is sufficient to provide a cushion against liquidity shocks. As a further alternative, one might impose narrow banking in the sense that banks are required to hold sufficient liquid funds so as to pay out in all contingencies. As shown in the next section, both these options turn out to be strictly worse than imposing minimum liquidity standards ex ante combined with lender-of-last-resort policy. They can even be inferior to the outcome of a mixed-strategy equilibrium with free-riding banks.

4. The Role of Equity and Narrow Banking

Let us now introduce equity requirements in the model; i.e., banks are required to hold some equity as a share of their assets. Instead of pure fixed deposit contracts, the banks now issue a mixture of deposit contracts and equity for attracting funds from the investors (Diamond and Rajan 2000, 2005, 2006). Equity can reduce the fragility by providing a cushion against negative shocks. This, however, comes at a cost, since it allows the bank manager to capture a rent. So the regulator needs to strike a balance between benefit and cost.

Being a renegotiable claim, in contrast to deposits, equity is subject to the hold-up problem; i.e., equity holders will only get a share ζ ($\zeta \in [0,1]$) of the surplus, the bank manager extracting the remaining part $1-\zeta$ as rent from his or her superior collection skills. Without changing the nature of the problem, in the following we simply assume that $\zeta = \frac{1}{2}$.

With $\zeta = \frac{1}{2}$, the bank manager and equity holders share the surplus over deposits equally. So the equity value of a bank not suffering from a run is $\frac{\gamma E[R_s]-d_0}{2}$ in state s with expected return $\gamma E[R_s]$ and deposit claims d_0. Assume that some equity requirement k is imposed—that is, the share of equity to bank assets is k, with

$$k = \frac{\frac{\gamma E[R_s]-d_0}{2}}{\frac{\gamma E[R_s]-d_0}{2} + d_0}.$$

Solving for d_0 gives the return to depositors as

$$d_0 = \frac{1-k}{1+k}\gamma E[R_s],$$

with equity holders receiving $\frac{k}{1+k}\gamma E[R_s]$. Thus investors providing funds both in the form of deposits and equity to the banks will receive the payoff $\frac{1}{1+k}\gamma E[R_s] < \gamma E[R_s]$ at $t = 1$.

In the absence of aggregate risk, introducing equity requirements is a pure cost, reducing the investors' payoff. Somewhat counterintuitively, without aggregate risk, equity requirements even reduce the share α invested in safe projects. The reason is that with equity financing bank managers get the rent $\frac{\gamma E[R]-d_0}{2}$, extracting part of the surplus over deposits from equity holders. Since the return at $t = 2$ is higher than at $t = 1$, bank managers are willing to consume late, so the amount of resources needed at $t = 1$ is lower in the presence of equity. Consequently, the share α will be reduced. Obviously, banks holding no equity provide more attractive conditions for investors, so equity could not survive. This seemingly counterintuitive result simply demonstrates that there is no role (or rather only a payoff-reducing role) for costly equity in the absence of aggregate risk.

The benefit of equity comes in when there is aggregate risk: Equity helps to absorb aggregate shocks and avoid the costly bank runs. In the simple two-state setup, equity holdings need to be just sufficient to cushion the bad state. With sufficient equity, the bank can choose $\alpha = \alpha_H$, profiting from the high return in the good state and still staying solvent in the bad state. In that case, it just needs to be able to pay out the fixed claims of investors, wiping out all equity.

With equity k and investment $\alpha = \alpha_H$, the total amount that can be pledged to investors providing funds both as depositors and equity holders is $\frac{1}{1+k}\gamma E[R_H]$ in the good state, with claims of depositors being $d_0 = \alpha_H R_1 + (1-\alpha_H)p_L R_2$ and return on equity $\frac{k}{1+k}\gamma E[R_H]$. In the bad state, a marginally solvent bank is able to pay out $d_0 = \alpha_H R_1 + (1-\alpha_H)p_L R_2$ to depositors. So the minimum k^* to prevent bank runs is determined by the condition

$$\frac{1-k^*}{1+k^*}\gamma E[R_H] = \alpha_H R_1 + (1-\alpha_H)p_L R_2,$$

and we solve to get

$$k^* = \frac{\gamma E[R_H] - d_0}{\gamma E[R_H] + d_0}. \tag{3}$$

Obviously, k^* is decreasing in p_L: The higher p_L, the lower the cushion k^* which is needed to stay solvent in the bad state.

Condition (3) determines the minimum equity requirement k^* a regulator needs to impose in order to eliminate the risk of costly bank runs. Setting k lower ($k < k^*$) would not help to prevent bank runs; setting k too high ($k > k^*$) would just raise the cost of holding equity without additional benefit. Thus from now on we can concentrate on the level k^* without loss of generality. In the following, we compare the investors' payoff in an economy subject to equity requirements with the payoff in the absence of any regulation (as derived in proposition 3) and then with the case of liquidity requirements, combined with the central bank acting as lender of last resort. Finally, we conclude with an analysis of narrow banking.

4.1 Equity Requirements versus Market Equilibrium

We first ask whether equity requirements can improve the investors' allocation in this economy, relative to the payoff they get in the market equilibrium we characterized in proposition 3. As shown in section 2, the investors' payoff depends on the probability π of the good state. If π is low enough, banks choose the safe strategy α_L. For high π, they pick α_H, with payoff increasing in π. In an intermediate range, free-riding banks drive down investors' return relative to what they could earn from investment in the safe strategy α_L. The overall payoff as a function of π is shown by the gray lines in figure 3. Let us call this function $\Pi(\pi)$. It seems natural to expect that equity requirements are superior at least for the intermediate range. As we will show, this intuition does not hold.

In figure 5A, the solid black lines show the investors' expected return $\Pi_e(\pi) = d_0 + \frac{\Pi}{2}\pi$ for the case of equity requirements. With equity requirements, the investors' expected return is uniformly increasing in π. The equity requirement k^* is chosen such that deposits will always be paid out fully, even in the bad state. Thus, the fixed deposit payment d_0 is independent of π. In contrast, the return on equity $\frac{\Pi}{2}$ is paid out only in the good state. The more likely the good state (the higher π), the higher is the expected return on equity. Its value is determined by

$$\frac{\Pi}{2} = \frac{\gamma E[R_H] - d_0}{2} = \frac{\gamma E[R_H] - \frac{1-k}{1+k}\gamma E[R_H]}{2} = \frac{k}{1+k}\gamma E[R_H].$$

Figure 5. Investors' Expected Return in the Market Equilibrium (Solid Gray Lines)/under Equity Requirements (Solid Black Lines)

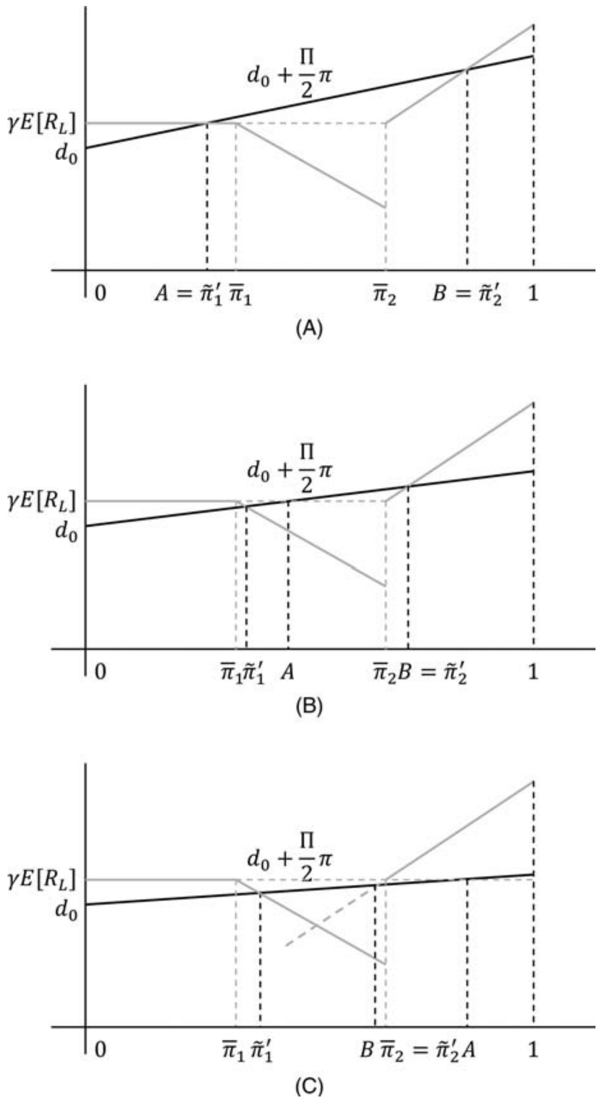

Note: In each of the three cases, the range where the equity regulation regime dominates the market equilibrium is denoted by the interval $(\tilde{\pi}_1', \tilde{\pi}_2')$.

Under what conditions will a banking system with equity requirements outperform the investors' return in the market equilibrium? Intuition suggests that relative performance depends on parameter values. As lemma 1 proves, equity requirements can never dominate the market outcome uniformly. It is straightforward to compare the investors' payoff under equity requirements with the market equilibrium with free-riding for the extreme values $\pi = 0$ and $\pi = 1$.

LEMMA 1. *The investors' expected return under the equity requirement is lower than the market equilibrium outcome when* $\pi = 0$ *or* $\pi = 1$.

Proof. See appendix 1.

The intuition of lemma 1 is straightforward: Since there is no uncertainty when $\pi = 0$ or $\pi = 1$, it is inferior to hold costly equities as explained above.

Figure 5A suggests, however, that equity requirements might uniformly improve the investor's expected return for the range of parameter values resulting in the mixed-strategy equilibrium with free-riding banks. Unfortunately, proposition 6 shows that this need not be the case. The equity regulation regime may even be sometimes dominated by the market equilibrium with free-riding.

PROPOSITION 6. *Imposing the equity requirement* k^* *may make investors better off than the mixed-strategy equilibrium with free-riding banks for some range of parameter values. But the costs of imposing equity requirements may be so high that the equity regulation regime may be dominated even by the market equilibrium with free-riding. There are three possible cases:*

(i) The equity regulation regime dominates the market equilibrium in the case of free-riding—that is, for $\overline{\pi}_1 \leq \pi \leq \overline{\pi}_2$.

(ii) In the range $\overline{\pi}_1 \leq \pi \leq \overline{\pi}_2$*, the equity regulation regime dominates for high values of* π*, whereas the market equilibrium with free-riding dominates for low values. In addition, the equity regulation regime dominates the market equilibrium for the low values of* π *in the range* $\overline{\pi}_2 \leq \pi \leq 1$.

(iii) In the range $\overline{\pi}_1 \leq \pi \leq \overline{\pi}_2$*, the equity regulation regime dominates for high values of* π*, whereas the market equilibrium*

with free-riding dominates for low values. In addition, the equity regulation regime is uniformly dominated by the market equilibrium in the range $\overline{\pi}_2 \leq \pi \leq 1$.

Proof. See appendix 1.

The three possible cases are characterized in figures 5A, B, and C, respectively. The quantitative conditions which separate these cases can be found in appendix 2. Numerical examples illustrating these cases are presented in appendix 3.

Proposition 6 says that the effectiveness of imposing equity requirements is dubious. Equity requirements may give investors a higher payoff than the mixed-strategy equilibrium with free-riding banks for all parameter values with mixed-strategy equilibrium $\overline{\pi}_1 \leq \pi \leq \overline{\pi}_2$. This case is captured as case (A) in figure 5A (and as proposition 6 (i)). Since free-riding partly destroys the value of assets held by prudent banks (forcing them to hold a riskier portfolio), it might seem that imposing equity requirements will always dominate the market equilibrium outcome with mixed strategies. But according to proposition 6, it is quite likely that equity requirements result in inferior payoffs for some range of parameter values (for example, when c is not very low and p_H is close to γ, i.e., the bank run cost is not very high), as shown in case (B) in figure 5B (when $A \in (\overline{\pi}_1, \overline{\pi}_2)$, as in proposition 6 (ii), equity requirements result in inferior payoffs to mixed-strategy equilibrium for $(\overline{\pi}_1, \tilde{\pi}_1')$ but superior to market equilibrium for some high values of $(\overline{\pi}_2, \tilde{\pi}_2'))$. It might be that imposing equity requirements makes investors even worse off, as in figure 5C, representing case (C) (when $A > \overline{\pi}_2$, as in proposition 6 (iii), equity requirements result in inferior payoffs to mixed-strategy equilibrium for $(\overline{\pi}_1, \tilde{\pi}_1')$ and inferior payoffs for all $\pi \in [\overline{\pi}_2, 1])$.

The intuition behind this result is that holding equity can be quite costly[3]; if so, it may be superior to accept the fact that systemic risk is a price to be paid for higher returns on average.

[3]The cost of holding equity comes from the fact that equity holders can only get the share ζ from the surplus in the good state. In this paper, we set $\zeta = 0.5$ for a simpler exposition. The investors' expected return increases when ζ gets higher and, ceteris paribus, outperforms the return in the market equilibrium for a wider range of π. In the limit, when $\zeta = 1$, issuing equity creates no hold-up problems and so does not incur costs. In that case, financing via equity would be able to implement the planner's solution.

Figure 6. Real Expected Return with Credible Liquidity Injections (Black Chain Line, for the Case of Figure 12), Compared with Investors' Expected Return in the Market Equilibrium (Solid Gray Lines)/under Equity Requirements (Solid Black Lines)

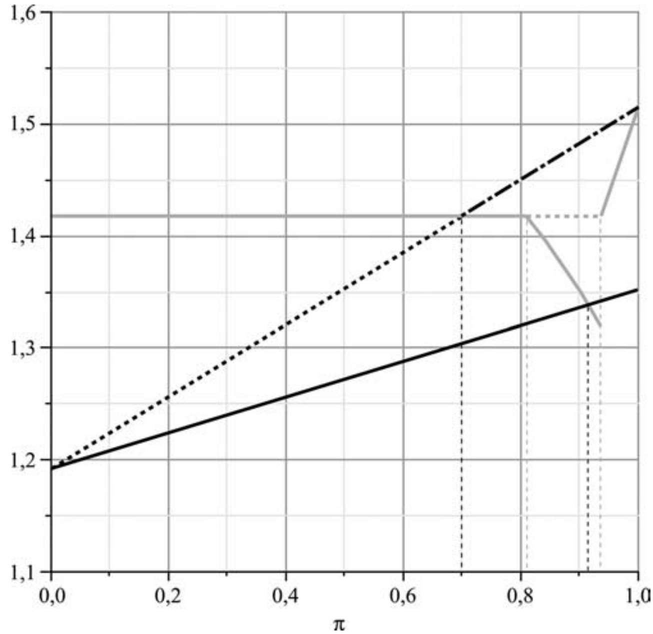

4.2 Equity Requirements versus Conditional Lender-of-Last-Resort Policy

As just shown, there is no clear ranking between market equilibrium without regulation and a regime with equity requirements. In contrast, the mix of ex ante liquidity requirements with ex post lender-of-last-resort policy is always dominating equity requirements. See figure 6. The reason is as follows: Consider that the banks are required to hold $\underline{\alpha} = \alpha_H$ when π is high. Then when p_H is revealed, the investors' real return is $\gamma E[R_H]$; and when p_L is revealed, the investors' real return is $\alpha_H R_1 + (1 - \alpha_H)p_L R_2$. Therefore the investors' overall expected return turns out to be

$$\Pi_m = \gamma E[R_H]\pi + (1 - \pi)[\alpha_H R_1 + (1 - \alpha_H)p_L R_2],$$

Figure 7. Real Expected Return with Narrow Banking Compared with Ex Ante Liquidity Regulation

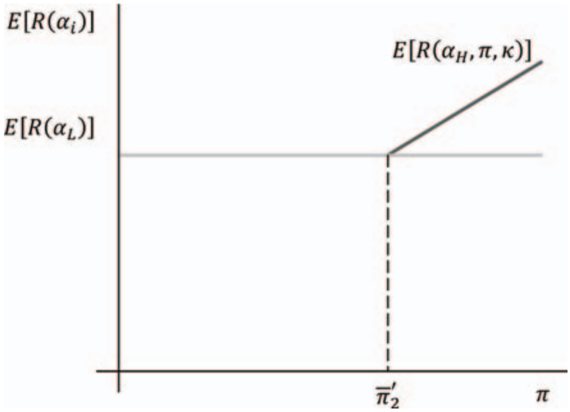

which is linear in π, as the chain line of figure 6 shows. Note that when $\pi = 1$, $\Pi_m = \gamma E[R_H] > d_0 + \frac{\Pi}{2}$; and when $\pi = 0$, $\Pi_m = \alpha_H R_1 + (1 - \alpha_H) p_L R_2 = d_0$. Therefore, Π_m line is above $d_0 + \frac{\Pi}{2}\pi$, $\forall \pi \in (0, 1]$; i.e., the mix of liquidity requirements with lender-of-last-resort policy is always dominating equity requirements when aggregate uncertainty exists.

4.3 Liquidity Requirements versus Narrow Banking

In times of crises, frequently there are calls to go back to narrow banking in order to avoid the risk of runs. Under narrow banking, institutions with deposits would be required to hold as assets only the most liquid instruments so as to be always able to meet any deposit withdrawal by selling their assets. Obviously, narrow banking can be extremely costly. In our model, banks would be required to hold sufficient liquid funds to pay out in all contingencies: $\alpha \geq \alpha_L$. As figure 7 illustrates, under narrow banking an investor's payoff (the solid gray line) can be much lower for high π compared with ex ante liquidity regulation combined with ex post lender-of-last-resort policy (the solid black line). Just as with equity requirements, narrow banking (imposing the requirement that banks hold sufficient equity so as to be able to pay out demand deposits in all states of the world) can be quite inferior: If the bad state is a rare probability event, it simply makes no sense to dispense with all the efficiency gains from investing in high-yielding illiquid assets despite its impact on systemic risk.

5. Conclusion

Traditionally, exposure to aggregate liquidity shocks has been modeled as an exogenous event. In this paper, we derive the aggregate share of liquid projects endogenously. It depends on the incentives of financial intermediaries to invest in risky, illiquid projects. This endogeneity allows us to capture the feedback between financial market regulation and incentives of private banks, determining the aggregate amount of liquidity available. As a consequence of limited liability, banks are encouraged to free-ride on liquidity provision. Relying on sufficient liquidity provided by the market, they are inclined to invest excessively in illiquid long-term projects.

Liquidity provision by central banks can help to prevent bank runs with inefficient early liquidation. However, the anticipation of unconditional liquidity provision results in overinvestment in risky activities (moral hazard), creating excessive exposure to systemic risk. We show that it is crucial for efficient lender-of-last-resort policy to impose ex ante minimum liquidity standards for banks. In addition, we analyze the impact of equity requirements. We show that it is even likely to be inferior to the outcome of a mixed-strategy equilibrium with free-riding banks. For similar reasons, imposing narrow banking (requiring banks to hold sufficient liquid funds to pay out in all contingencies) turns out to be strictly inferior relative to the combination of liquidity requirements with lender-of-last-resort policy.

In modern economies, a significant part of intermediation is provided by the shadow banking sector. These institutions (like hedge funds and investment banks) are not financed via deposits, but they are highly leveraged. Incentives to dance (to free-ride on liquidity provision) seem to be even stronger for the shadow banking industry. So imposing liquidity requirements only for the banking sector will not be sufficient to cope with free-riding. In future work, we plan to analyze incentives for leveraged institutions within our framework.

Appendix 1. Proofs

Proof of Proposition 1

In the absence of aggregate risk, given p_s $(s \in \{H, L\})$, the social planner maximizes the investors' return by setting α_s such that

$$\alpha_s = \arg \max_{\alpha_s \in [0,1]} \gamma \left\{ \alpha_s R_1 + (1 - \alpha_s) \left[p_s R_2 + \frac{(1 - p_s) R_2}{r} \right] \right\}, \quad \text{with } r_s \geq 1. \tag{4}$$

Solve to get $\alpha_s = \frac{\gamma - p_s}{\gamma - p_s + (1-\gamma)\frac{R_1}{R_2}}$, with $r_s = 1$.

In the presence of aggregate risk, to find the social planner's optimal α which may depend on π, one just has to find the α that maximizes the investors' return for each $\pi \in [0, 1]$.

That the gross interest rate offered to the entrepreneurs at $t = 1$ is no less than 1 implies that for any given α, the investors' expected payoff is

$$E[R(\alpha)] = \pi \min\{\alpha R_1 + (1 - \alpha) p_H R_2, \gamma[\alpha R_1 + (1 - \alpha) R_2]\} + (1 - \pi) \min\{\alpha R_1 + (1 - \alpha) p_L R_2, \gamma[\alpha R_1 + (1 - \alpha) R_2]\},$$

which is linear in π. Then it is easy to depict $E[R(\alpha)]$ as a function of π, when $\alpha = \alpha_H$ or α_L, as figure 8 shows. These two lines intersect at $\overline{\pi}_2' = \frac{\gamma E[R_L] - \kappa}{\gamma E[R_H] - \kappa}$. Note that $E[R(\alpha_H)] = \gamma E[R_H] > \gamma E[R_L]$ when $\pi = 1$, and $E[R(\alpha_H)] = \kappa < \gamma E[R_L]$ when $\pi = 0$.

For any $\alpha \in (\alpha_L, 1]$, $E[R(\alpha)] = \gamma[\alpha R_1 + (1 - \alpha) R_2] < \gamma E[R_L]$, as shown by the dotted gray lines in figure 8. For any $\alpha \in [0, \alpha_H)$, $E[R(\alpha)] = \pi[\alpha R_1 + (1 - \alpha) p_H R_2] + (1 - \pi)[\alpha R_1 + (1 - \alpha) p_L R_2]$. Note that $E[R(\alpha)] < \kappa$ when $\pi = 0$ and $E[R(\alpha)] < \gamma E[R_H]$ when $\pi = 1$, as shown by the dotted black lines in figure 8.

For any $\alpha \in (\alpha_H, \alpha_L)$, $E[R(\alpha)] = \pi\gamma[\alpha R_1 + (1 - \alpha) R_2] + (1 - \pi)[\alpha R_1 + (1 - \alpha) p_L R_2]$. Denote $\alpha R_1 + (1 - \alpha) R_2$ by $E[R_\alpha]$, and

Figure 8. The Investors' Expected Return for any $\boldsymbol{\alpha \in [0, 1]}$

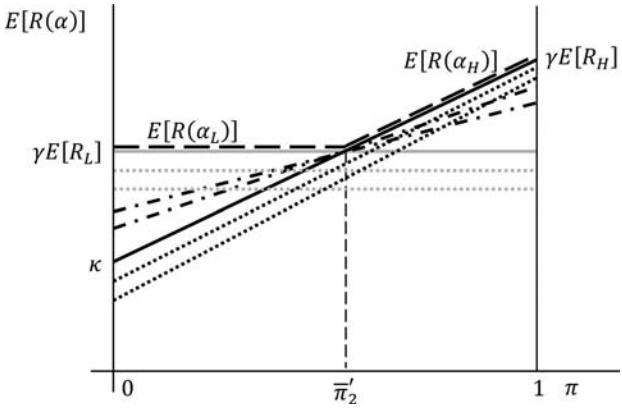

Notes: The solid gray line denotes $E[R(\alpha_L)]$, the black line denotes $E[R(\alpha_H)]$, the dotted gray lines denote those $E[R(\alpha)]$ with $\alpha \in (\alpha_L, 1)]$, the dotted black lines denote those $E[R(\alpha)]$ with $\alpha \in (0, \alpha_H]$, and the chain lines denote those $E[R(\alpha)]$ with $\alpha \in (\alpha_H, \alpha_L]$.

$\alpha R_1 + (1-\alpha)p_L R_2$ by κ'. Note that $\kappa < E[R(\alpha)] < \gamma E[R_L]$ when $\pi = 0$ and $\gamma E[R_L] < E[R(\alpha)] < \gamma E[R_H]$ when $\pi = 1$. Such $E[R(\alpha)]$ are depicted as the chain lines in figure 8.

Suppose that the intersection between $E[R(\alpha)]$ and $E[R(\alpha_L)]$ is $\overline{\pi}_2'' = \frac{\gamma E[R_L]-\kappa'}{\gamma E[R_\alpha]-\kappa'}$. To determine the value of $\overline{\pi}_2''$, note that $\overline{\pi}_2'' \gtrless \overline{\pi}_2'$ only if $\frac{\gamma E[R_L]-\kappa'}{\gamma E[R_\alpha]-\kappa'} \gtrless \frac{\gamma E[R_L]-\kappa}{\gamma E[R_H]-\kappa}$. This is equivalent to

$$\gamma E[R_L](\gamma E[R_H] - \gamma E[R_\alpha]) + (\gamma E[R_\alpha] - \gamma E[R_L])\kappa + (\gamma E[R_L] - \gamma E[R_H])\kappa' \gtrless 0. \tag{5}$$

Using the fact that $\gamma E[R_s] = \alpha_s R_1 + (1-\alpha_s)p_s R_2$ ($s \in \{H, L\}$), the left-hand side of the inequality (5) can be written as

$$\gamma(R_1 - p_L R_2)\{E[R_H](\alpha_L - \alpha) - E[R_\alpha](\alpha_L - \alpha_H) + E[R_L](\alpha - \alpha_H)\}.$$

Further, since $\alpha \in (\alpha_H, \alpha_L)$, we can replace α with the linear combination of α_H and α_L, $\alpha = \omega\alpha_H + (1-\omega)\alpha_L$ with $\omega \in (0, 1)$. It is easily seen that

$$\gamma(R_1 - p_L R_2)\{E[R_H](\alpha_L - \alpha) - E[R_\alpha](\alpha_L - \alpha_H) \\ + E[R_L](\alpha - \alpha_H)\} = 0,$$

which implies that $\overline{\pi}_2'' = \overline{\pi}_2'$.

Combining all the cases, figure 8 shows the investors' expected return for any $\alpha \in [0,1]$. The social planner's optimal solution is given by the frontier of the investors' expected return (as shown by the dashed black lines in figure 8), which is a state-contingent strategy depending on the probability π: The planner invests the share α_H in liquid projects as long as $\overline{\pi}_2' \leq \pi \leq 1$, and the share α_L in liquid projects as long as $0 \leq \pi < \overline{\pi}_2'$.

Proof of Proposition 2

To show that the optimal allocation of the central planner's problem is supported by the market equilibrium, one has to show that (i) the allocation is feasible in the market economy, and (ii) it is not profitable to unilaterally deviate from such allocation.

In the planner's economy, the central planner picks up the optimal α_s as equation (4) suggests, and transfers the maximized return to the investors. This coincides with equations (1) and (2), implying that claim (i) holds.

To show that claim (ii) holds, suppose an arbitrary bank i deviates from such allocation by choosing $\alpha_i \neq \alpha_s$ for a given $s \in \{H, L\}$:

(i) If $\alpha_i < \alpha_s$, by the market clearing condition, the liquidity market interest rate r' at $t = 1$ is now determined by

$$r'\{(1-\gamma)[\alpha_i R_1 + (1-\alpha_i)p_s R_2] \\ + (N-1)(1-\gamma)[\alpha_s R_1 + (1-\alpha_s)p_s R_2]\} \\ = \gamma(1-\alpha_i)(1-p_s)R_2 + (N-1)\gamma(1-\alpha_s)(1-p_s)R_2.$$

Comparing with the condition in the central planner's problem in which $r = 1$,

$$r(1-\gamma)[\alpha_s R_1 + (1-\alpha_s)p_s R_2] = \gamma(1-\alpha_s)(1-p_s)R_2,$$

one can see that $r' > 1$. For the non-deviators, the depositors' return becomes

$$\gamma\left\{\alpha_s R_1 + (1-\alpha_s)\left[p_s R_2 + \frac{(1-p_s)R_2}{r'}\right]\right\} < d_0.$$

Knowing that the non-deviators will not be able to meet the contracted d_0 at $t = 1$, the depositors will only deposit at bank i at $t = 0$. If so, the deposit return that bank i can offer is at maximum

$$d_{0i} = \alpha_i R_1 + (1-\alpha_i)p_s R_2 < d_0,$$

implying that the deviator gets worse off.

(ii) If $\alpha_i > \alpha_s$, the aggregate liquidity supply at $t = 1$ exceeds the the aggregate liquidity demand because

$$\begin{aligned}(1-\gamma)[\alpha_i R_1 + (1-\alpha_i)p_s R_2] + (N-1)(1-\gamma)[\alpha_s R_1 + (1-\alpha_s)p_s R_2] \\ > N(1-\gamma)[\alpha_s R_1 + (1-\alpha_s)p_s R_2] \\ = N\gamma(1-\alpha_s)(1-p_s)R_2.\end{aligned}$$

Therefore, the liquidity market interest rate remains at $r = 1$ and the non-deviators are able to meet d_0. However, the deposit return that bank i can offer is

$$d_{0i} = \gamma[\alpha_i R_1 + (1-\alpha_i)R_2] < \gamma[\alpha_s R_1 + (1-\alpha_s)R_2] = d_0,$$

implying that the deviator will not get any deposit at $t = 0$ and is hence worse off.

Therefore, the planner's optimal allocation is indeed supported by the market equilibrium.

Proof of Proposition 3

The mixed-strategy equilibrium, proposition 3(i)–3(iii)(c), is characterized as proposition 2 of Cao and Illing (2008). By choosing to hold a share of safe assets—call it α_s^*—a prudent bank should have equal return at both states, $d_0^s = d_0^s(p_H) = d_0^s(p_L)$, i.e.,

$$\gamma\left[\alpha_s^* R_1 + (1-\alpha_s^*)p_H R_2 + \frac{(1-\alpha_s^*)(1-p_H)R_2}{r_H}\right]$$
$$= \gamma\left[\alpha_s^* R_1 + (1-\alpha_s^*)p_L R_2 + \frac{(1-\alpha_s^*)(1-p_L)R_2}{r_L}\right].$$

With some simple algebra, this is equivalent to

$$\frac{1}{r_H} = \frac{1-p_L}{1-p_H}\frac{1}{r_L} - \frac{p_H-p_L}{1-p_H}.$$

Plot $\frac{1}{r_H}$ as a function of $\frac{1}{r_L}$, as the solid black line in figure 9 shows.

The slope $\frac{1-p_L}{1-p_H} > 1$ and intercept $-\frac{p_H-p_L}{1-p_H} < 0$, and the line goes through $(1,1)$. But $r_H = r_L = 1$ cannot be the equilibrium outcome here, because α_L is the dominant strategy in this case and subject to deviation. So whenever $r_H > 1$ (suppose $\frac{1}{r_H} = A$ in the graph), there must be $r_H > r_L > 1$ (because $\frac{1}{r_H} < \frac{1}{r_L} = B < 1$).

At p_L, given that $r_L > 1$, the prudent bank's return is equal to $d_0^s = \kappa(\alpha_s^*(p_L, r_L)) < \kappa(\alpha_L)$, since the latter maximizes the bank's

Figure 9. Higher Interest Rates in the Mixed-Strategy Equilibrium

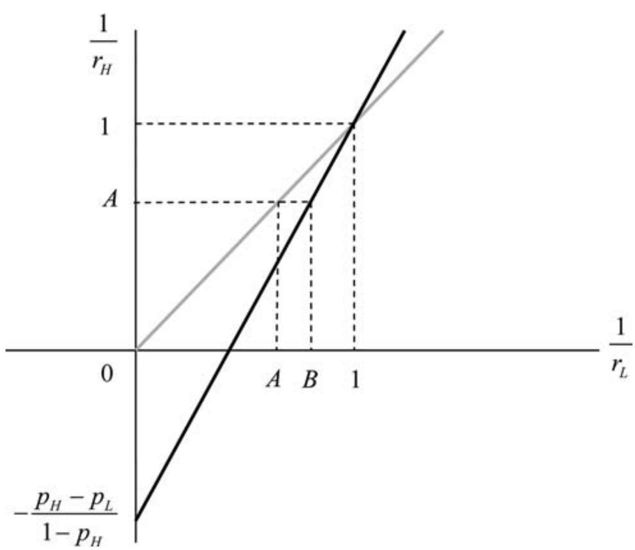

expected return with $r^* = 1$ by lemma 2 of Cao and Illing (2008). Therefore in the mixed-strategy equilibrium, investors are worse off than if all banks would coordinate on the prudent strategy α_L.

Proof of Proposition 4

The central bank's optimal policy is to restore the constrained efficiency, as stated in proposition 1. Therefore, the optimal liquidity requirement, which is captured by $\underline{\alpha}$, should be exactly the same as α_H (α_L) for high (low) π. So is it with r^M.

In addition, any bank that observes $\underline{\alpha}$ will get bailed out whenever necessary. This only happens when a bank follows $\underline{\alpha} = \alpha_H$ but p_L is revealed. In this case, the investors will get $\kappa = \alpha_H R_1 + (1 - \alpha_H)p_L R_2$ real return plus $d_0 - \kappa$ fiat money if they do not run on the bank at $t = \frac{1}{2}$. In contrast, the investors will only get the liquidated value $c < 1 < \kappa$ as the real return if they run on the bank. Of course, they will wait instead of run.

Proof of Proposition 5

Suppose that a representative bank chooses to be prudent with $\alpha_i = \underline{\alpha}$, and promises a nominal deposit contract $d_0^i = \gamma[\underline{\alpha}R_1 + (1-\underline{\alpha})R_2]$ in order to maximize its investors' return. Then when the bad state with high liquidity needs is realized, the central bank has to inject enough liquidity into the market to keep the interest rate at $r = 1$ in order to ensure bank i's survival. However, given $r = 1$, a naughty bank j can always profit from setting $\alpha_j = 0$, promising the nominal return $d_0^j = \gamma R_2 > d_0^i$ to its investors. Thus, surely the banks prefer to play naughty.

For those parameter values such that $\pi p_H R_2 + (1-\pi)p_L R_2 < 1$, there exists no equilibrium with liquidity injection. The reason is the following:

(i) Any symmetric strategic profile cannot be in equilibrium, because
   - (a) if there is no trade under such strategic profile—i.e., α is so small that the real return is less than 1—one bank can deviate by setting $\alpha = 1$ and trading with investors;
   - (b) if there is trade under such strategic profile—i.e., $\alpha > 0$ for all the banks—then one bank can deviate by setting

$\alpha = 0$ and getting higher nominal return than the other banks.

(ii) Any asymmetric strategic profile, or profile of mixed strategies, cannot be in equilibrium, because
   (a) if there is no trade under such strategic profile, then the argument of (i)(a) applies here;
   (b) if there is trade under such strategic profile, then one bank can deviate by choosing a pure strategy, $\alpha = 0$, and get better off—there is no reason to mix with the other dominated strategies.

Proof of Lemma 1

When $\pi = 0$,

$$d_0 + \frac{\Pi}{2} \cdot 0 = \alpha_H R_1 + (1 - \alpha_H) p_L R_2 < \alpha_L R_1 + (1 - \alpha_L) p_L R_2$$
$$= \gamma E[R_L].$$

When $\pi = 1$,

$$d_0 + \frac{\Pi}{2} = \frac{\alpha_H R_1 + (1 - \alpha_H) p_L R_2 + \alpha_H R_1 + (1 - \alpha_H) p_H R_2}{2}$$
$$< \alpha_H R_1 + (1 - \alpha_H) p_H R_2$$
$$= \gamma E[R_H].$$

Proof of Proposition 6

As lemma 1 shows, the investors' expected return with equity requirements $d_0 + \frac{\Pi}{2}\pi$ is a linear increasing function of π, starting from $d_0 < \gamma E[R_L]$ when $\pi = 0$ and ending with $d_0 + \frac{\Pi}{2} < \gamma E[R_H]$ when $\pi = 1$. Whether imposing equity requirements improves investors' expected return (such "improved" region is denoted by the interval $(\tilde{\pi}_1', \tilde{\pi}_2')$ in figure 5) depends on the intersection between $d_0 + \frac{\Pi}{2}\pi$ and $\gamma E[R_L]$, denoted by A as in figure 5. Generically, there are three cases concerning the relative positions of $\Pi(\pi)$ and $\Pi_e(\pi)$:

(i) As figure 5A shows, when $A \in (0, \overline{\pi}_1]$, $\Pi_e(\pi)$ is higher than $\Pi(\pi)$ for $\pi \in [\overline{\pi}_1, \overline{\pi}_2]$.

(ii) As figure 5B shows, when $A \in (\overline{\pi}_1, \overline{\pi}_2)$, $\Pi_e(\pi)$ is only higher than $\Pi(\pi)$ for part of $\pi \in [\overline{\pi}_1, \overline{\pi}_2]$. In addition, $\Pi_e(\pi)$ is higher for part of $\pi \in (\overline{\pi}_2, 1]$.

(iii) As figure 5C shows, when $A \geq \overline{\pi}_2$, $\Pi_e(\pi)$ is only higher than $\Pi(\pi)$ for part of $\pi \in [\overline{\pi}_1, \overline{\pi}_2]$. In addition, $\Pi_e(\pi)$ is no higher for all $\pi \in (\overline{\pi}_2, 1]$.

Appendix 2. Effectiveness Conditions for Equity Requirements

To economize the notations, define the investors' expected return function in the market equilibrium as follows:

DEFINITION 1. *Define a representative investor's expected return function without equity requirements as* $\Pi(\pi)$, *such that*

$$\Pi(\pi, \cdot) = \begin{cases} \gamma E[R_L], & \text{if} \quad \pi \in [0, \overline{\pi}_1]; \\ \alpha_s^* R_1 + \big(1 - \alpha_s^*\big) p_L R_2, & \text{if} \quad \pi \in (\overline{\pi}_1, \overline{\pi}_2); \\ \gamma E[R_H]\pi + (1 - \pi)c, & \text{if} \quad \pi \in [\overline{\pi}_2, 1] \end{cases}$$

and his or her expected return function under equity requirements as $\Pi_e(\pi)$, *as well as the set* S *in which the investor's payoff is improved under equity requirements, such that*

$$S := \{\hat{\pi} | \Pi_e(\hat{\pi}) \geq \Pi(\hat{\pi})\}.$$

In the following, we are interested in the cases, captured in the set S (denoted by the area $(\tilde{\pi}_1', \tilde{\pi}_2')$ in figure 5), where the banking system with equity requirements outperforms that in the market equilibrium.

Denote the intersection of $\Pi_e(\pi) = d_0 + \frac{\Pi}{2}\pi$ and $\gamma E[R_L]$ by A, which is equal to (see the proof below for detail)

$$A = \frac{2(R_1 - p_L R_2)}{(1 - \gamma)R_1 + (\gamma - p_L)R_2},$$

as well as the intersection of $\Pi_e(\pi) = d_0 + \frac{\Pi}{2}\pi$ and $\gamma E[R_H]\pi + (1 - \pi)c$ by B, which is equal to (see the proof below for detail)

$$B = \frac{2[(1-\gamma)(cR_1 - p_L R_1 R_2) + (\gamma - p_H)(cR_2 - R_1 R_2)]}{2(1-\gamma)cR_1 + 2(\gamma - p_H)cR_2 + [\gamma(p_H - 1) - (\gamma - p_H) - (1-\gamma)p_L]R_1 R_2}.$$

Then proposition 7 characterizes the improvement in investors' payoff achieved by introducing equity requirements.

PROPOSITION 7. *Given the equity requirement k imposed by the regulator,*

(i) When $A \in (0, \overline{\pi}_1]$, i.e.,

$$(2\gamma R_2 - \gamma E[R_H] - d_0)(\gamma E[R_L] - d_0) \\ + (2\gamma E[R_L] - \gamma E[R_H] - d_0)(d_0 - c) \leq 0,$$

then $S = [A, B] \supseteq [\overline{\pi}_1, \overline{\pi}_2]$;

(ii) When $A \in (\overline{\pi}_1, \overline{\pi}_2)$, i.e.,

$$(2\gamma R_2 - \gamma E[R_H] - d_0)(\gamma E[R_L] - d_0) \\ + (2\gamma E[R_L] - \gamma E[R_H] - d_0)(d_0 - c) > 0,$$

and

$$\gamma(E[R_H] - E[R_L])(d_0 - c) \geq (\gamma E[R_H] - c)(\gamma E[R_L] - d_0),$$

then $S = [\tilde{\pi}, B]$ in which $\tilde{\pi} \in (\overline{\pi}_1, \overline{\pi}_2)$ and $S \bigcap [\overline{\pi}_1, \overline{\pi}_2] = [\tilde{\pi}, \overline{\pi}_2]$;

(iii) When $A \geq \overline{\pi}_2$, i.e.,

$$2(\gamma E[R_L] - d_0)(\gamma E[R_H] - c) \geq (\gamma E[R_H] - d_0)(\gamma E[R_L] - c),$$

then $S \subseteq [\tilde{\pi}, B]$ in which $\tilde{\pi} \in (\overline{\pi}_1, \overline{\pi}_2)$ and $S \bigcap [\overline{\pi}_1, \overline{\pi}_2] = [\tilde{\pi}, \overline{\pi}_2]$.

Proof. The intersection A takes the value of π, such that

$$\gamma E[R_L] = d_0 + \frac{\Pi}{2}\pi.$$

Solve to get

$$A = \frac{2(\gamma E[R_L] - d_0)}{\gamma E[R_H] - d_0} = \frac{2(R_1 - p_L R_2)}{(1-\gamma)R_1 + (\gamma - p_L)R_2}.$$

The intersection B takes the value of π, such that

$$\gamma E[R_H]\pi + (1-\pi)c = d_0 + \frac{\Pi}{2}\pi.$$

Solve to get

$$B = \frac{d_0 - c}{\frac{\gamma E[R_H]+d_0}{2} - c}$$
$$= \frac{2[(1-\gamma)(cR_1 - p_L R_1 R_2) + (\gamma - p_H)(cR_2 - R_1 R_2)]}{2(1-\gamma)cR_1 + 2(\gamma - p_H)cR_2 + [\gamma(p_H - 1) - (\gamma - p_H) - (1-\gamma)p_L]R_1 R_2}.$$

Then the set S can be determined in each case:

(i) As figure 5A shows, when $A \in (0, \overline{\pi}_1]$,

$$\frac{2(\gamma E[R_L] - d_0)}{\gamma E[R_H] - d_0} \leq \overline{\pi}_1 = \frac{\gamma E[R_L] - c}{\gamma R_2 - c}.$$

Rearrange to get

$$(2\gamma R_2 - \gamma E[R_H] - d_0)(\gamma E[R_L] - d_0) + (2\gamma E[R_L] - \gamma E[R_H] - d_0)(d_0 - c) \quad \leq 0.$$

Since $\Pi_e(\pi)$ is strictly increasing in π, then

$$\begin{aligned} &\Pi_e(\pi)|_{\pi=B} > \Pi_e(\pi)|_{\pi=A} \geq \gamma E[R_L]|_{\pi=\overline{\pi}_1} \\ &= [\gamma E[R_H]\pi + (1-\pi)c]|_{\pi=\overline{\pi}_2} \\ &\geq \Pi(\pi)|_{\pi\in[\overline{\pi}_1,\overline{\pi}_2]}, \end{aligned}$$

which implies $S = [A, B] \supseteq [\overline{\pi}_1, \overline{\pi}_2]$;

(ii) As figure 5B shows, when $A \in (\overline{\pi}_1, \overline{\pi}_2]$,

$$\overline{\pi}_1 = \frac{\gamma E[R_L] - c}{\gamma R_2 - c} < \frac{2(\gamma E[R_L] - d_0)}{\gamma E[R_H] - d_0}.$$

Rearrange to get

$$(2\gamma R_2 - \gamma E[R_H] - d_0)(\gamma E[R_L] - d_0) + (2\gamma E[R_L] - \gamma E[R_H] - d_0)(d_0 - c) \quad > 0.$$

What's more, in this case $B \in [\overline{\pi}_2, 1]$, and this is equivalent to

$$\frac{\gamma E[R_L] - c}{\gamma E[R_H] - c} = \overline{\pi}_2 < \frac{d_0 - c}{\frac{\gamma E[R_H] + d_0}{2} - c}.$$

Rearrange to get

$$\gamma(E[R_H] - E[R_L])(d_0 - c) \geq (\gamma E[R_H] - c)(\gamma E[R_L] - d_0).$$

Similarly,

$$\begin{aligned}\Pi_e(\pi)|_{\pi \leq A} &\leq \gamma E[R_L]|_{\pi=\overline{\pi}_1} = [\gamma E[R_H]\pi + (1-\pi)c]|_{\pi=\overline{\pi}_2} \\ &\leq \Pi(\pi)|_{\pi \in [\overline{\pi}_2, B]} \\ &\leq \Pi_e(\pi)|_{\pi \geq B},\end{aligned}$$

which implies $S = [\tilde{\pi}, B]$ in which $\tilde{\pi} \in (\overline{\pi}_1, \overline{\pi}_2]$ and $S \bigcap [\overline{\pi}_1, \overline{\pi}_2] = [\tilde{\pi}, \overline{\pi}_2]$;

(iii) As figure 5C shows, when $A \geq \overline{\pi}_2$,

$$\overline{\pi}_2 = \frac{\gamma E[R_L] - c}{\gamma E[R_H] - c} \leq \frac{2(\gamma E[R_L] - d_0)}{\gamma E[R_H] - d_0}.$$

Rearrange to get

$$2(\gamma E[R_L] - d_0)(\gamma E[R_H] - c) \geq (\gamma E[R_H] - d_0)(\gamma E[R_L] - c).$$

Similarly,

$$\begin{aligned}&\Pi_e(\pi)|_{\pi \leq B} < \Pi_e(\pi)|_{\pi \geq A} \leq \gamma E[R_L]|_{\pi=\overline{\pi}_1} \\ &= [\gamma E[R_H]\pi + (1-\pi)c]|_{\pi=\overline{\pi}_2},\end{aligned}$$

which implies $S \subseteq [\tilde{\pi}, B]$ in which $\tilde{\pi} \in (\overline{\pi}_1, \overline{\pi}_2]$ and $S \bigcap [\overline{\pi}_1, \overline{\pi}_2] = [\tilde{\pi}, \overline{\pi}_2]$.

Appendix 3. Numerical Examples

The following figures present numerical illustrations representing the three different cases.

Figure 10. Investors' Expected Return in the Market Equilibrium (Solid Gray Lines)/under Equity Requirements (Solid Black Lines), with $p_H = 0.3$, $p_L = 0.25$, $\gamma = 0.6$, $R_1 = 1.8$, $R_2 = 5.5$, and $c = 0.9$

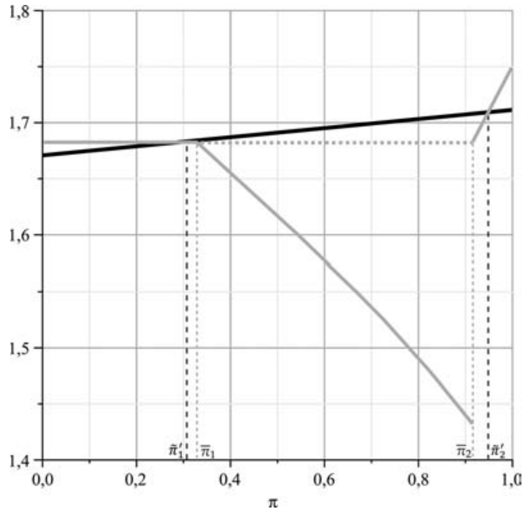

Figure 11. Investors' Expected Return in the Market Equilibrium (Solid Gray Lines)/under Equity Requirements (Solid Black Lines), with $p_H = 0.4$, $p_L = 0.3$, $\gamma = 0.6$, $R_1 = 2$, $R_2 = 4$, and $c = 0.8$

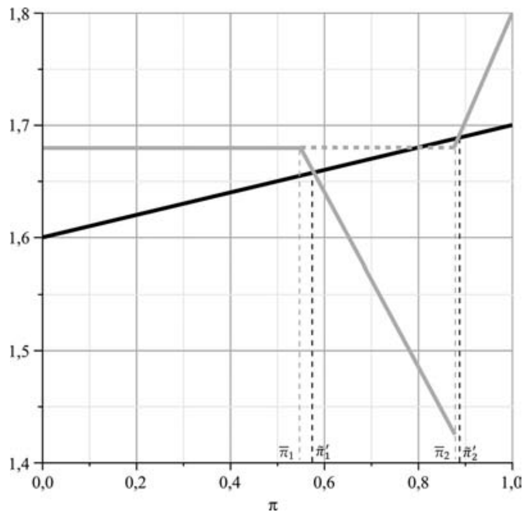

Figure 12. Investors' Expected Return in the Market Equilibrium (Solid Gray Lines)/under Equity Requirements (Solid Black Lines), with $p_H = 0.5$, $p_L = 0.25$, $\gamma = 0.7$, $R_1 = 1.8$, $R_2 = 2.5$, and $c = 0$

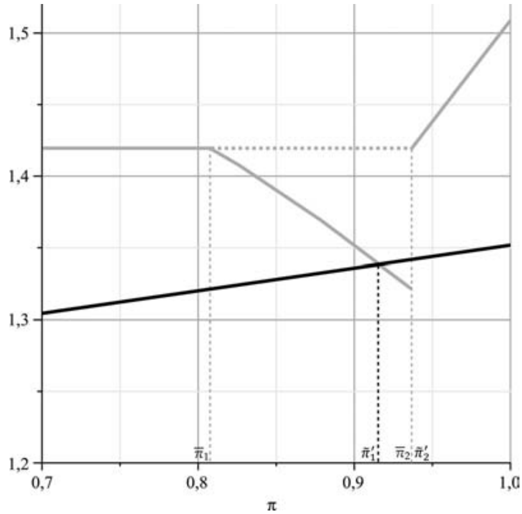

References

Allen, F., E. Carletti, and D. Gale. 2010. "Money, Financial Stability, and Efficiency." Working paper.

Allen, F., and D. Gale. 1998. "Optimal Financial Crises." *Journal of Finance* 53 (4): 1245–84.

———. 2004. "Financial Intermediaries and Markets." *Econometrica* 72 (4): 1023–61.

———. 2007. *Understanding Financial Crisis.* Clarendon Lectures, New York: Oxford University Press.

Bagehot, W. 1873. "A General View of Lombard Street." Reprinted in *Financial Crises, Contagion, and the Lender of Last Resort: A Reader,* ed. C. Goodhart and G. Illing. New York: Oxford University Press, 2002.

Buiter, W. H. 2007. "What Did You Do in the Open Market Today, Daddy?" Blog entry. Available at http://blogs.ft.com/maverecon/2007/12/.

Buiter, W. H., and A. C. Sibert. 2007. "The Central Bank as Market Maker of Last Resort." Blog entry. Available at http://blogs.ft.com/maverecon/2007/08/the-central-banhtml/.

Cao, J., and G. Illing. 2008. "Liquidity Shortages and Monetary Policy." CESifo Working Paper No. 2210. Available at SSRN: http://ssrn.com/abstract=1090825.

Diamond, D. W., and P. H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91 (3): 401–19.

Diamond, D. W., and R. G. Rajan. 2000. "A Theory of Bank Capital." *Journal of Finance* 55 (6): 2431–65.

———. 2001. "Liquidity Risk, Liquidity Creation and Financial Fragility: A Theory of Banking." *Journal of Political Economy* 109 (2): 287–327.

———. 2005. "Liquidity Shortage and Banking Crises." *Journal of Finance* 60 (2): 615–47.

———. 2006. "Money in a Theory of Banking." *American Economic Review* 96 (1): 30–53.

Farhi, E., and J. Tirole. 2009. "Collective Moral Hazard, Maturity Mismatch and Systemic Bailouts." NBER Working Paper No. 15138. Available at http://www.nber.org/papers/w15138.

Goodfriend, M., and R. King. 1988. "Financial Deregulation, Monetary Policy, and Central Banking." In *Restructuring Banking and Financial Services in America*, AEI Studies No. 481, ed. W. S. Haraf and R. M. Kushmeider. Lanham, MD: UPA.

Hart, O., and J. Moore. 1994. "A Theory of Debt Based on the Inalienability of Human Capital." *Quarterly Journal of Economics* 109 (4): 841–79.

Holmström, B., and J. Tirole. 1997. "Financial Intermediation, Loanable Funds, and the Real Sector." *Quarterly Journal of Economics* 112 (3): 663–91.

———. 1998. "Private and Public Supply of Liquidity." *Journal of Political Economy* 106 (1): 1–40.

Rochet, J.-C. 2004. "Macroeconomic Shocks and Banking Supervision." *Journal of Financial Stability* 1 (1): 93–110.