Why the Publication of Socially Harmful Information May Be Socially Desirable*

Volker Hahn
CER-ETH
Center of Economic Research
at ETH Zurich
ZUE D13
8092 Zurich, Switzerland
vhahn@ethz.ch

First Version: May 2009
This Version: October 2010

Abstract

We propose a signaling model in which the central bank and firms receive information on cost-push shocks independently of each other. If the firms are rather unlikely to receive information directly, central banks should remain silent about their own private information. If, however, firms are sufficiently likely to be informed, it is socially desirable for the central bank to reveal its own private information. By doing so, the central bank eliminates the distortions stemming from the signaling incentives under opacity. The possibility of the central bank withholding information on a discretionary basis does not limit the usefulness of a transparency requirement for ameliorating social welfare. Moreover, our model may provide a rationale for the recent trend towards more transparency in monetary policy.

Keywords: signaling games, transparency, monetary policy, central banks, communication.

JEL: D82, D83, E58.

*I would like to thank Johannes Becker, Hans Gersbach, Oliver Grimm, Stephan Imhof, Maik Schneider, and participants of the EEA annual meeting in Glasgow, 2010, for many valuable comments and suggestions.

1 Introduction

During the last two decades, central banks' communication practices have changed dramatically. While traditionally central banks were wrapped in mystery and withheld information about their policies, their assessment of the economy, details of decision-making and the goals of monetary policy, they have gradually become substantially more open. In 1987, the then chairman of the Federal Reserve Board, Alan Greenspan, took pride in being secretive: "Since I've become a central banker," he noted, "I've learned to mumble with great incoherence."[1] Nowadays such a statement would be unthinkable. For example, the present chairman Ben Bernanke called the "increased openness" of monetary-policy makers a "welcome development" in 2007.[2]

In this paper, we examine whether this development is socially beneficial. For this purpose, we present a simple model, populated by a central bank that receives private information about cost-push shocks and a continuum of firms that may receive information from the central bank or through other sources. As shown by Angeletos and Pavan (2007) in their Corollary 9, private agents' information about socially inefficient sources of business-cycle fluctuations reduces welfare. The fact that cost-push shocks represent such socially inefficient sources of fluctuations therefore suggests that central banks should always aim at keeping information about these shocks secret. In this paper, we derive the seemingly paradoxical finding that central-bank transparency with regard to cost-push shocks may nevertheless be socially desirable.

Our finding can be explained in the following way. First, central-bank transparency has the effect that the mere publication of information about socially undesirable shocks may trigger socially undesirable fluctuations in situations where firms would be unaware of the shocks otherwise. Second, a more subtle effect arises. Even if the central bank does not publish its private information, the firms may infer some information by observing the monetary policy conducted by the central bank. Consequently, the central bank not only has to consider the direct impact of its action on price setters, it

[1] Wall Street Journal, September 22, 1987.

[2] "Federal Reserve Communications", speech delivered at the Cato Institute 25th Annual Monetary Conference, Washington, D.C. November 14, 2007.

also has to take into account the effect that its action has on their expectations about the shock. Due to this effect, opacity reduces the central bank's flexibility in stabilizing shocks.

On balance, we show that the aggregate consequences of both effects for welfare depend on the probability of firms receiving independent information. First, if this probability is low, opacity is socially desirable because it reduces the detrimental economic effects caused by cost-push shocks. The central bank can achieve this by pursuing a passive policy, thereby safeguarding the secrecy of its information. Second, for a sufficiently high probability of firms' receiving independent information, the central bank would not remain inactive if it were opaque. As the central bank will thus reveal its information anyway, transparency is socially desirable because it removes the restraint imposed by the link between the central bank's actions and the firms' expectations under opacity.

We also consider the possibility that, under a formal transparency requirement, central banks may have the discretionary power to withhold information. If this is actually beneficial to central banks, a transparency requirement will be largely ineffective. The phenomenon that central banks can circumvent transparency requirements is not just a theoretical possibility. When there were attempts to force the Fed to publish the minutes of its committee's meetings, the Fed tried to abolish the minutes altogether (see Lewis (1991)). In our model, we show that central banks that are required to be transparent with respect to cost-push shocks will not withhold their information discretionarily if transparency is socially desirable. Equivalently, whenever transparency is desirable ex ante, it will also be desirable ex post after the central bank has observed the shock.

While an important reason for the trend towards openness in monetary policy is that central banks have given in to outside pressure (by Buiter (1999), for example), our paper identifies an additional mechanism that may have contributed to this trend. With the constant progress of information technologies, the precision of private agents' direct information about economic developments is likely to have been improving over time. This development may have made transparency in monetary policy more attractive.

The assumption that the firms' information has been improving is consistent with the empirical evidence presented in D'Agostino and Whelan (2008). This evidence suggests that the superiority of the Fed's forecasts, which was identified in Romer and Romer (2000), has vanished recently.[3]

Our paper is organized as follows: In the next section, we review the related literature. We outline the model and conduct some first steps of our analysis in sections 3 and 4. In sections 5 and 6, we derive the equilibria under transparency and opacity, respectively. We compare welfare under both transparency regimes in Section 7. The circumstances under which a transparency requirement is effective are examined in Section 8. The robustness of our results is considered in Section 9. Section 10 concludes.

2 Related Literature

Our paper contributes to two strands of the literature. First, it is part of the literature on signaling games, which goes back to Spence (1973). In monetary economics, signaling games have been studied by Gersbach and Hahn (2007, 2009), Sibert (2002, 2003, 2009), and Vickers (1986). In our paper, the public will attempt to infer the central bank's private information about economic shocks from monetary-policy action if the central bank keeps this information secret. This is consistent with the empirical finding in Romer and Romer (2000) that contractionary monetary policy leads to increases in inflation expectation. This counter-intuitive result can be explained if we assume that contractionary monetary policy signals private information of the central bank about shocks that drive up future inflation.

In our paper, we consider not only separating equilibria like most papers on signaling games in monetary policy, but study also pooling and semi-separating equilibria. This enables us to identify the new effect that pooling and semi-separating equilibria enable the central bank to ensure complete or partial secrecy of its information.

Second, our paper complements the general literature on transparency in monetary policy as surveyed by Geraats (2002), Hahn (2002), and Blinder et al. (2008). This

[3]However, this effect could also be explained by increases in central-bank transparency.

literature considers the economic effects of central-bank communication as well as the consequences that the publication of private central-bank information has for welfare.[4] Cukierman and Meltzer (1986) show that central banks may prefer some degree of ambiguity about monetary control in order to be able to surprise private agents at a time when this is most valuable to them. This framework has been modified to allow for normative analysis (see Lewis (1991)) and for an explicit distinction of control-error variance and the degree of transparency (Faust and Svensson (2001)). Applying a New Keynesian specification of the Phillips curve, Jensen studies the desirability of transparency with regard to the central bank's control error (Jensen (2002)) and private information about cost-push shocks (Jensen (2000)). In contrast to Jensen (2000), firms' price-setting is affected by the firms' expectations about the shock rather than the actual shock realization in the present paper.

Our paper is also related to Baeriswyl and Cornand (2007, 2010) and Walsh (2007). These contributions study the dual role of the central bank's instrument as a stabilization tool and a public signal of the central bank's private information. The authors assume that the central bank commits to a linear rule that minimizes unconditional losses and focus on the public's signal-extraction problem when the central bank's instrument responds to two different shocks. By contrast, we make the assumption that the central bank chooses its policy on a discretionary basis, after observing its private information. We concentrate on the non-linear distortion in the central bank's response to a single private signal under opacity.[5] To the best of our knowledge, this paper is also the first to examine the circumstances under which central banks may use their discretionary power to withhold evidence even if they are operating under a transparency regime.

[4] In a much-cited paper, Morris and Shin (2002) show that transparency may be socially harmful if agents find it individually optimal to coordinate their actions and if this coordination is not socially desirable *per se*. However, Svensson (2006) convincingly argues that the range of parameters for which this result holds is unlikely to be relevant in practice. Hellwig (2005) highlights that transparency may lead to socially desirable coordination in price-setting.

[5] Because of our assumption of discretionary central-bank policy, we obtain results that are different from those found in previous papers. For example, Baeriswyl and Cornand (2007) find that transparency is detrimental with respect to markup shocks. The main point of our paper is that transparency may be beneficial even in this case.

3 Model

We consider a one-shot signaling game with a central bank (the sender) and a multitude of firms (receivers). For the sake of brevity, we abstain from presenting microeconomic foundations to our model, which can be found in Adam (2007), for example.

In line with Eq. (3) in Adam (2007), each firm's optimal price, p^*, is given by

$$p^* = p + \alpha y + \varepsilon', \tag{1}$$

where p is the aggregate price level, y is the (log) output gap, α a positive parameter, and ε' represents a cost-push shock.[6]

On theoretical grounds, cost-push shocks can be justified by variations in markups. Markup shocks can be modeled by a stochastic sales tax on all goods, where revenues are used to finance lump-sum transfers to the agents.[7] Markup shocks can also be motivated by changes in the intensity of competition or in the aggressiveness of wage bargainers.[8] A general feature of markup shocks is that it is individually optimal to firms to react to them. However, this response leads to socially undesirable fluctuations that cannot be stabilized by the central bank perfectly. In this sense, information about cost-push shocks is socially harmful in the hands of price-setters. This is a special case of the more general insight provided by Angeletos and Pavan (2007) in their Corollary 9 that private and public information about sources of inefficient business-cycle fluctuation is welfare-reducing.

In order to keep the model tractable, we assume that only four realizations of the shock are possible.[9] The shock realizations are $-e'_L$, $+e'_L$, $-e'_H$, $+e'_H$ with $0 < e'_L < e'_H$ (where L stands for low and H for high). The prior probabilities are ρ_L for $-e'_L$ and $+e'_L$ and ρ_H for $-e'_H$ and $+e'_H$ $(2\rho_H + 2\rho_L = 1)$. Thus the shock distribution is symmetric. The restriction to a discrete set of possible shocks enables us not only to derive analytical

[6] We normalize (log) natural output to zero and thus use the terms "output" and "output gap" interchangeably.

[7] For a discussion of markup shocks see Ball et al. (2005), among others. Compare also the extensive discussion in Woodford (2002), pp. 44-45.

[8] In Adam (2007), cost-push shocks are labeled real demand shocks.

[9] The distortions identified in this paper would also materialize for other shock distributions, as we will explain in Section 9.

results for separating equilibria but also to study the existence of pooling equilibria and semi-separating equilibria, where some types of central banks pool, while others choose a policy that perfectly reveals their type. By contrast, most analyses of signaling games in monetary economics are restricted to separating equilibria (see, for example, Sibert (2009)).

The central bank observes the shock with probability p_{CB} ($0 \leq p_{CB} \leq 1$). With the complementary probability, the central bank obtains no private information. Consequently, there exist five types of central banks. Four types correspond to the different possible shocks. The fifth type, denoted by 0, is a central bank that has not observed the shock. The set of possible types is thus $\mathcal{T} := \{-H, -L, 0, +L, +H\}$. Under transparency, the central bank's type is published. It is kept secret under opacity.

The central bank chooses its instrument m (log money growth), which affects output via a quantity equation:

$$y = m - p \tag{2}$$

The central bank's loss function, which also represents social losses, is given by

$$\mathcal{L} = \frac{1}{1+a}p^2 + \frac{a}{1+a}y^2, \tag{3}$$

where $a \geq 0$ is a parameter that measures the importance of the output target (this loss function has been derived by Adam (2007) in his Appendix A.2, for example; see also Woodford (2002)).[10]

Adopting the notion of sticky information that has been introduced by Mankiw and Reis (2002), we assume that only a fraction λ ($0 < \lambda < 1$) of firms use up-to-date information. These firms, which we henceforth refer to as "attentive," observe monetary-policy decisions and also the information that the central bank publishes if it is operating under a transparency regime.[11] In addition, these firms jointly observe ε' with

[10] Compared to the more standard formulation $\mathcal{L} = p^2 + ay^2$, we have normalized losses by the factor $\frac{1}{1+a}$. Obviously, this does not affect our findings; however, it will simplify the analysis.

[11] In principle, one could also consider the case where the central bank's instrument is kept secret. For example, the Federal Reserve did not make its policy directive public immediately after the board meetings (see Goodfriend (1986)). However, the central bank's instrument, usually a short-term interest rate, can easily be observed. Therefore it is implausible that monetary-policy makers can keep their choices of instrument secret for any considerable length of time.

probability p_F ($0 \leq p_F \leq 1$), the event of them observing the shock being independent of whether the central bank observes the shock. The remainder of firms are inattentive and do not obtain any information about the shock or the central bank's choice of monetary policy.

The sequence of events is as follows:

1. Nature draws the shock ε'.

2. The central bank becomes informed about ε' with probability p_{CB}.

3. Under opacity, the central bank's information is kept private. Under transparency, it is published.

4. The central bank chooses its instrument m.

5. A fraction λ of firms, i.e. the attentive firms, obtain precise information on ε' with probability p_F. With the complementary probability, these firms remain ignorant of the shock. Moreover, the attentive firms observe m and any information the central bank might have published. The remaining firms do not update their information.

6. All firms choose the prices of their outputs.

In line with other papers that study signaling games in monetary economics, we do not study an infinite-horizon model (see, for example, Sibert (2009)). The focus on one period keeps our framework analytically tractable. We note that our results are driven by the distortions created by the central bank's signaling incentives under opacity. These distortions would also occur in a variant of our model with multiple periods.

We will derive perfect Bayesian Nash equilibria for two scenarios: a transparency scenario in which the central bank publishes its observation of the shock if it has observed the shock and an opacity scenario in which the central bank keeps this information secret.

4 Preliminary Steps

We begin our analysis with a few preliminary steps that enable us to present the central bank's losses in a convenient form. As inattentive firms' expectations about ε' are zero, the (log) prices they choose are also zero. As a result, the price level is given by $p = \lambda \mathbb{E}_F[p^*]$, where $\mathbb{E}_F$ is the expectations operator with respect to the attentive firms' information set. The equilibrium price level as a function of m can now be obtained by inserting $y = m - p$ into (1) and applying $\mathbb{E}_F$:

$$p = \lambda \left(p + \alpha(m - p) + \mathbb{E}_F[\varepsilon']\right) \tag{4}$$

Solving for p yields

$$p = \frac{\lambda}{1 - \lambda(1 - \alpha)} \left(\alpha m + \mathbb{E}_F[\varepsilon']\right). \tag{5}$$

It will be useful to introduce normalized values for the shock $\varepsilon := \lambda/(1 - \lambda(1 - \alpha))\varepsilon'$ and its realizations $e_\tau := [\lambda/(1 - \lambda(1 - \alpha))]\, e'_\tau \ \forall \tau \in \mathcal{T} \setminus \{0\}$. With the help of $\sigma := (\lambda\alpha)/(1 - \lambda(1 - \alpha))$, p can now be rewritten as

$$p = \sigma m + \mathbb{E}_F[\varepsilon]. \tag{6}$$

It is crucial to note that $\mathbb{E}_F[\varepsilon]$ is determined by the attentive firms' direct observation of the shock if they have in fact observed the shock. If they have not observed the shock directly, $\mathbb{E}_F[\varepsilon]$ depends on the central bank's information under transparency and on the central bank's choice of m under opacity.

Inserting (2) and (6) into (3) yields

$$\mathcal{L}(m, \mathbb{E}_F[\varepsilon]) = \frac{1}{1 + a}(\sigma m + \mathbb{E}_F[\varepsilon])^2 + \frac{a}{1 + a}((1 - \sigma)m - \mathbb{E}_F[\varepsilon])^2. \tag{7}$$

Importantly, the central bank could always achieve zero losses by choosing $m = 0$ if the attentive firms' expectations concerning the cost-push shock were zero. By contrast, the central bank can never achieve zero losses when firms' expectations are different from zero. Thus information about cost-push shocks is socially harmful.

5 Transparency

In this section we focus on the transparency scenario. In the following, we derive the optimal policy chosen by the different types of central bank. For types in $\mathcal{T} \setminus \{0\}$, we obtain $\mathbb{E}_F[\varepsilon] = \varepsilon$. It is straightforward to check that (7) can be rewritten in the following way:

$$\mathcal{L}(m, \mathbb{E}_F[\varepsilon]) = \frac{a}{(1+a)(\sigma^2 + a(1-\sigma)^2)}(\mathbb{E}_F[\varepsilon])^2 + \frac{\sigma^2 + a(1-\sigma)^2}{1+a}(m - m^T_{\mathbb{E}_F[\varepsilon]})^2, \quad (8)$$

where

$$m^T_{\mathbb{E}_F[\varepsilon]} := \frac{a - \sigma(1+a)}{\sigma^2 + a(1-\sigma)^2}\mathbb{E}_F[\varepsilon]. \quad (9)$$

Variable $m^T_{\mathbb{E}_F[\varepsilon]}$ can be interpreted as the optimal value of m, conditional on the fact that the attentive firms' expectations about ε are given by $\mathbb{E}_F[\varepsilon]$. With slight abuse of notation we will sometimes write m^T_τ for $m^T_{\mathbb{E}_F[\varepsilon]}$ evaluated at $\mathbb{E}_F[\varepsilon] = e_\tau$ ($\tau \in \mathcal{T}$). Then m^T_τ represents the optimal choices of central banks $\tau \in \mathcal{T} \setminus \{0\}$ under transparency.

For $a - \sigma(1+a) > 0$, $m^T_{\mathbb{E}_F[\varepsilon]}$ is a strictly monotonically increasing function of $\mathbb{E}_F[\varepsilon]$. For $a - \sigma(1+a) < 0$, it is strictly monotonically decreasing. This observation is important, as we will draw an analogy under opacity and impose monotonicity as a restriction on the equilibria. To simplify the exposition, we exclude the knife-edge case $a - \sigma(1+a) = 0$ for the remainder of the paper.

It then remains to derive the optimal policy of an uninformed type $\tau = 0$. $\mathbb{E}_F[\varepsilon]$ may take five different values from this central bank's perspective, namely $-e_H$, $-e_L$, 0, e_L, e_H, depending on whether the attentive firms receive information about the shock and, if so, which realization they observe. Hence an uninformed central bank chooses m to minimize expected losses

$$\mathcal{L}_0 := p_F \left[\rho_L \mathcal{L}(m, -e_L) + \rho_L \mathcal{L}(m, +e_L) + \rho_H \mathcal{L}(m, -e_H) + \rho_H \mathcal{L}(m, +e_H)\right] + (1-p_F)\mathcal{L}(m, 0).$$

Importantly, $\mathcal{L}(m, -e_L) + \mathcal{L}(m, +e_L)$, $\mathcal{L}(m, -e_H) + \mathcal{L}(m, +e_H)$, and $\mathcal{L}(m, 0)$ are quadratic functions of m with minima at $m = 0$. As a consequence, the optimal policy of an uninformed central bank under transparency is given by $m^T_0 := 0$.

We summarize our observations in the following proposition:

Proposition 1

Under transparency, a unique equilibrium exists. Each central bank of type $\tau \in \mathcal{T}$ chooses m_τ^T.

In this equilibrium, each central bank chooses m so as to optimally trade off the effect of the shock on output and prices. Because the central bank makes its private information public, it does not have to care about its choice of m affecting the firms' estimate of the shock.

6 Opacity

Under opacity, the attentive firms do not receive the central bank's information directly. However, upon observing the central bank's choice of money growth, they may update their estimate of the central bank's type. Under opacity, the model thus corresponds to a signaling game.

With probability p_F, attentive firms learn the correct realization of ε because they receive information independently of the central bank. With probability $1 - p_F$ they obtain no independent information and attempt to infer the central bank's information from the central bank's choice of money growth m. We introduce $f(m)$ with $-e_H \leq f(m) \leq e_H \ \forall m$ to denote the attentive firms' expectations about ε, given that they have not observed the shock.

We focus on perfect Bayesian Nash equilibria in pure strategies that satisfy two additional, plausible assumptions on $f(m)$. In these perfect Bayesian Nash equilibria, the firms' beliefs about the central bank's type and thus $f(m)$, in particular, will be consistent with the central bank's equilibrium strategy.

First, we impose a monotonicity requirement on $f(m)$, as will be detailed in the following. Imposing monotonicity is intuitive, given that under transparency the central bank's choice of m is a monotonic function of its estimate about the shock. Under transparency, the equilibrium value of m is an increasing function of $\mathbb{E}_{CB}[\varepsilon]$ for $a - \sigma(1 + a) > 0$ and a decreasing function for $a - \sigma(1 + a) < 0$. Hence we assume that attentive firms' expectations under opacity are a monotonic function of m. In particular, we postulate that $f(m)$ is weakly increasing for $a-\sigma(1+a)>0$ and weakly decreasing for $a-\sigma(1+a)<0$.

Second, we assume that $f(m)$ is an odd function, i.e. $f(m)=-f(-m)\ \forall m$. This is plausible because of the model's linear-quadratic nature. Under transparency, for example, the central bank's optimal choice of m is also an odd function of the central bank's estimate about the shock.

These assumptions have several important implications. First, firms expect the shock to be zero if they have not observed it and the central bank has chosen $m=0$. Formally, this can be stated as $f(0)=0$. Second, and consequently, a central bank of type 0 will choose $m=0$, as can be verified easily. Third, the equilibrium choices of all types $\mathcal{T}$ are a weakly monotonic function of the central bank's estimate of the shock. Formally, this implies $m^O_{-H}\leq m^O_{-L}\leq m^O_0=0\leq m^O_{+L}\leq m^O_{+H}$ for $a-\sigma(1+a)>0$ and $m^O_{-H}\geq m^O_{-L}\geq m^O_0=0\geq m^O_{+L}\geq m^O_{+H}$ for $a-\sigma(1+a)<0$, using m^O_τ to denote type τ's equilibrium choice of m under opacity ($\tau\in\mathcal{T}$).[12]

In the following it will prove useful to introduce a critical value of p_F, denoted by p^*_F, as follows

$$p^*_F := \frac{a}{(\sigma^2+a(1-\sigma)^2)(1+a)}. \tag{10}$$

It is straightforward to derive

$$1-p^*_F = \frac{(a-\sigma(1+a))^2}{(\sigma^2+a(1-\sigma)^2)(1+a)} > 0.$$

Hence one can conclude $0\leq p^*_F<1$.[13]

We are now in a position to describe the equilibria under opacity. In Appendix A we prove an important proposition:

[12]The third consequence of our assumptions can be explained as follows. Suppose, for example, $a-\sigma(1+a)>0$. Then $f(m)\geq 0\ \forall m>0$ because $f(0)=0$ and $f(m)$ is weakly monotonically increasing. A central bank that has observed a positive shock would never choose a negative money growth rate $m<0$ because $-m>0$ would yield lower losses, which is readily verified with the help of (7). As a consequence, H and L choose positive values of m. Analogously, $-H$ and $-L$ choose negative values of m. Monotonicity of $f(m)$ then requires $m_H\geq m_L$ (otherwise the firms' beliefs would be incorrect). In a similar vein, $m_{-H}\leq m_{-L}$ follows from the monotonicity of $f(m)$.

[13]Recall that we have excluded the knife-edge case $a-\sigma(1+a)=0$.

Proposition 2

For $p_F < p_F^*$, a unique[14] *equilibrium exists. In this equilibrium, all types of central banks $\tau \in \mathcal{T}$ choose $m = 0$. If the attentive firms have not received direct information about ε, their expectations about the shock are $f(0) = 0$. For $a - \sigma(1 + a) > 0$, $f(m) = e_H$ $\forall m > 0$ and $f(m) = -e_H$ $\forall m < 0$. For $a - \sigma(1 + a) < 0$, $f(m) = -e_H$ $\forall m > 0$ and $f(m) = +e_H$ $\forall m < 0$.*

The proof is given in Appendix A. Intuitively, if the chances of the firms receiving information directly are rather low, it is profitable for the central bank to remain completely passive. As the firms are unlikely to learn about the shock, the expected losses incurred by not stabilizing the shock are low. Importantly, by not responding to its own private information, the central bank can prevent the firms from inferring this information.

For $p_F > p_F^*$, no unique equilibrium exists in general. In the following we will characterize several different equilibria for this case. With the help of

$$\hat{p}_F := \frac{e_H + e_L}{e_H + (2p_F^* - 1)e_L} p_F^*, \tag{11}$$

it is possible to describe the circumstances under which the same outcome as under transparency can prevail under opacity:

Proposition 3

If and only if $p_F \geq \hat{p}_F$, there is an equilibrium under opacity in which all types of central bankers $\tau \in \mathcal{T}$ choose the money growth rates they would find optimal under transparency (m_τ^T).

For the proof, see Appendix B. Intuitively, for high values of p_F the attentive firms are likely to be informed about the shock directly. As a consequence, it is optimal for the central bank to behave in the same manner as under transparency.

We note that $\hat{p}_F < 1$. Hence, the range of values of p_F for which the fully separating equilibria described in Proposition 3 exist is always non-empty. Additionally, we note

[14] To be more precise, the equilibrium is unique in the sense that no additional equilibrium exists in which the equilibrium choices for the five central bank types $\mathcal{T}$ are different. However, additional equilibrium with different out-of-equilibrium beliefs exist.

that $\hat{p}_F > p_F^*$. Consequently, for the interval $p_F^* < p_F < \hat{p}_F$ neither pooling equilibria, which are described in Proposition 2, nor separating equilibria with the same choices as under transparency exist. Intuitively, separating equilibria with the same choices as under transparency do not exist, as there would be strong incentives for central banks of type H to mimic L. By mimicking the L-type, type H can reduce the firms' expectations about the shock, which leads to lower losses if the firms do not receive information independently. However, if H could successfully mimic L, this would be costly to L as the firms might mistake it for H. This, in turn, would lead to high losses due to the firms' beliefs that the shock is very large. Thus type L tends to choose an m farther away from m_H^T in order to make mimicking more costly for H.

One example of such behavior is given in the following proposition, proven in Appendix C:

Proposition 4

There is a critical value for p_F, denoted by $\widetilde{p}_F$, with $p_F^* < \widetilde{p}_F < 1$ such that the following semi-separating equilibrium exists under opacity for $p_F \in [p_F^*, \widetilde{p}_F]$. Central banks of types $\tau \in \{-L, 0, +L\}$ choose $m = 0$. Central banks of types $\tau \in \{-H, +H\}$ choose m_τ^T. If attentive firms have not received direct information about ε, their expectations about the shock are $f(0) = 0$. For $a - \sigma(1 + a) > 0$, $f(m) = e_H$ $\forall m > 0$ and $f(m) = -e_H$ $\forall m < 0$. For $a - \sigma(1 + a) < 0$, $f(m) = -e_H$ $\forall m > 0$ and $f(m) = +e_H$ $\forall m < 0$.

We have shown that, for sufficiently large values of p_F ($p_F \geq \hat{p}_F$), separating equilibria exist where all central-bank types display the same behavior under opacity as under transparency. Moreover, for sufficiently small values of p_F ($p_F \leq p_F^*$), pooling equilibria exist. For somewhat larger p_F ($p_F^* \leq p_F \leq \widetilde{p}_F$), semi-separating equilibria occur in which central banks of types $-L$ and L mimic the behavior of 0. We note that $\widetilde{p}_F < \hat{p}_F$ cannot be ruled out, as can be readily verified. Thus it now remains to describe equilibria for the interval $p_F \in]\widetilde{p}_F; \hat{p}_F[$. This gap is filled by the following proposition:

Proposition 5

Suppose $\widetilde{p}_F < \hat{p}_F$. For all $p_F \in]\widetilde{p}_F; \hat{p}_F[$, values $\underline{\phi}$ and $\overline{\phi}$ with $0 < \underline{\phi} < \overline{\phi} < 1$ exist such that for all $\phi \in [\underline{\phi}; \overline{\phi}]$ separating equilibria exist under opacity that satisfy the following properties: Central banks of types $-H$ and $+H$ choose m^T_{-H} and m^T_H, respectively. Central banks of types $-L$ and $+L$ choose ϕm^T_{-L} and ϕm^T_L, respectively. type 0 chooses $m^T_0 = 0$.

The proof is given in Appendix D. These equilibria are particularly interesting as they represent fully separating equilibria where the behavior of types $-L$ and L is distorted by the factor ϕ over and against the equilibria under transparency. This distortion is the result of the incentives of types $-H$ and H to mimic the behavior of the types with moderate shock realizations, i.e. $-L$ and L. As successful imitation may increase the firms' shock estimate, types $-L$ and L choose the distorted money growth rates ϕm^T_{-L} and ϕm^T_L respectively, which makes mimicking less attractive for $-H$ and H.

To sum up, we have demonstrated that a perfect Bayesian Nash equilibrium satisfying our additional assumptions on $f(m)$ always exists. In each of these equilibria, the firms' beliefs $f(m)$ are consistent with the central bank's equilibrium strategy.

7 Comparison

In this section, we compare the central bank's losses and thus also social losses under transparency with the losses under opacity. The following proposition, proved in Appendix E, contains the major finding of this paper:

Proposition 6

For $p_F < p^*_F$, transparency is strictly inferior to opacity. For $p^*_F < p_F < \hat{p}_F$, transparency is strictly superior. For $p_F \geq \hat{p}_F$, transparency is weakly superior.[15]

We stress that Proposition 6 does not only hold with regard to the equilibria characterized in the previous section. It holds for all perfect Bayesian Nash equilibria satisfying our additional assumptions about $f(m)$.

[15] For $p_F \geq \hat{p}_F$, transparency and opacity lead to equivalent results with respect to welfare if the equilibria specified in Proposition 3 materialize. Transparency is strictly superior for all other equilibria.

Proposition 6 has the implication that whether transparency is desirable depends on the quality of the firms' direct information. If attentive firms are unlikely to be well-informed, transparency is detrimental. If there is a high probability of their being well-informed, central-bank transparency is desirable.

The intuition for this finding is as follows. If the central bank publishes its private information, it provides the attentive firms with information that may be unknown to them. As it is individually optimal for firms, albeit socially harmful, to respond to the shock, publishing information is costly to society. On the other hand, transparency eliminates the signaling costs the central bank incurs if the money growth rate it would like to choose under transparency were to signal the wrong information under opacity.

For low-quality information available to firms (and corresponding low levels of p_F), the costs incurred by transparency outweigh the benefits. Loosely speaking, it is better to remain inactive in this case and to speculate that firms will not discover the shock realization. By contrast, if the firms' information is high in quality, the firms will be probably informed anyway. By publishing its private information the central bank can avoid the signaling costs.

8 Feasibility of Transparency

In this section, we focus on the question whether transparency is feasible. Even if the central bank is required to be transparent, it is plausible that it can always withhold evidence. If this was actually profitable to the central bank for some shock realizations, a transparency requirement would be ineffective even in the case where it would be socially desirable.

In order to examine the feasibility of transparency, we extend the transparency scenario in the following way. Whenever the central bank observes the shock, it can choose between publishing this information and asserting that it does not know about a shock. However, it cannot forge information by claiming that it observed a shock that it did not observe. We will call transparency feasible if a perfect Bayesian Nash equilibrium of this modified game exists in which the central bank never withholds its private information.

In Appendix 8, we show

Proposition 7

Transparency is feasible if $p_F \geq p_F^*$.

A comparison with Proposition 7 reveals that transparency is feasible exactly in those circumstances in which it is socially desirable. Thus the possibility of the central bank withholding private information does not limit the usefulness of transparency from a social perspective.

9 Discussion

Here we discuss some issues related to the robustness of our findings. In particular, we focus on the specification of shocks, different types of shocks, the quality of the central bank's information, and the additional restrictions on equilibrium we have introduced under opacity.

Specification of shocks In this paper we have focused on four different shock realizations. This number is sufficiently high to identify the important signaling incentives in our framework and at the same time low enough to permit analytical results. If we considered only one possible realization of a positive and a negative shock (as opposed to the two in our model), we would ignore the crucial incentive of type H to mimic type L, which leads to the distortions under opacity driving our results regarding welfare. By contrast, if we considered more possible realizations of positive and negative shocks, the signaling incentives and thus the distortions would remain, but the analysis would be substantially more complex. In particular, with a continuum of potential shock realizations it is possible to show that pooling equilibria exist under opacity for small values of p_F and fully separating equilibria occur for large values of p_F, which is in line with the analysis in this paper.[16]

[16] A detailed analysis is available upon request.

Other types of shocks In our paper we deliberately focus on cost-push shocks because we intend to demonstrate that even with these shocks transparency can be socially desirable. We could introduce demand shocks into our framework, but transparency regarding these shocks would never be socially harmful. Under opacity, separating equilibria exist that would perfectly reveal the central bank's information. Then transparency and opacity would be equivalent with respect to welfare. Moreover, additional equilibria may exist under opacity, which would definitely entail lower welfare levels (for a detailed analysis, see Hahn (2009)). Consequently, transparency would be desirable from a welfare point of view.

Quality of the central bank's information Interestingly, the quality of the central bank's information, which is associated with parameter p_{CB} in our model, is irrelevant for the relative performance of transparency and opacity. Consequently, our findings extend to the case where central banks are always informed about cost-push shocks.

Restrictions on equilibrium In Section 6, where we analyze the opacity scenario, we have introduced two important restrictions on the equilibria under opacity, namely that $f(m)$ is monotonic and odd. Relaxing these assumptions might allow for additional equilibria. Although a complete characterization of all additional equilibrium candidates is beyond the scope of this paper, it is plausible that these equilibria would lead to higher losses under opacity. For example, an equilibrium where type L chooses negative values of m under opacity despite $m_L^T > 0$ is likely to be less desirable than equilibria satisfying our restrictions. Hence relaxing the restrictions on equilibria might make transparency more attractive over and against opacity.

10 Conclusions

In this paper, we have addressed the question whether central banks should publish information on sources of socially inefficient business-cycle fluctuations. Although, at first sight, it seems that the central bank should withhold evidence on these disturbances, we have shown that transparency may be socially desirable even in this case. If the probability of firms receiving information independently is sufficiently high, transparency eliminates the signaling incentives of different types of central banks and hence, in turn, the policy distortions prevalent under opacity. Our analysis has also highlighted that transparency is always feasible if it is socially desirable, as withholding information on a discretionary basis is not beneficial to the central bank in this case.

Our model can also be used to rationalize the current trend towards transparency in monetary policy to some extent. As improvements in information technologies plausibly raise the probability of economic agents receiving information on the economy independently of the central bank, it may be increasingly important for central banks to become more open about their assessment of the economy. We certainly do not want to argue that the mechanism identified in this paper is the only explanation for the increased transparency of central banks, but it may have contributed to this ongoing development. At any rate, our analysis suggests that the increased openness of monetary-policy makers is in fact a welcome development.

A Proof of Proposition 2

A.1 Existence

To show that the proposed equilibrium exists, we have to demonstrate that there is no profitable deviation for all types $\tau \in \mathcal{T}$. Before we show this, we note that central-bank losses can be written in a compact manner with the help of p_F^* (see (10)). Using (8) and (9), we obtain

$$\mathcal{L}(m_e^T, \mathbb{E}_F[\varepsilon]) = p_F^*(\mathbb{E}_F[\varepsilon])^2 + (1 - p_F^*)(e - \mathbb{E}_F[\varepsilon])^2. \tag{12}$$

This expression gives the losses the central bank incurs if attentive firms believe the shock to be $\mathbb{E}_F[\varepsilon]$ and if the central bank chooses the money growth rate m_e^T, which is the choice that would be optimal under transparency given that attentive firms would believe the shock to be e.

Deviations for 0 It is obvious that there is no profitable deviation for 0, as $m = 0$ is its preferred choice under transparency as well and any other choice would imply that the public believes a large shock has occurred, which would increase losses further. It thus remains to show whether profitable deviations exist for the other types. For simplicity, we focus on the case $a - \sigma(1 + a) > 0$. In this case, $f(m) = e_H \; \forall m > 0$ and $f(m) = -e_H \; \forall m < 0$ hold. The case with $a - \sigma(1 + a) < 0$ is completely analogous and is therefore omitted.

Deviations for H and -H Now we focus on possible deviations for H. In equilibrium, type H's losses are

$$p_F \mathcal{L}(0, \mathbb{E}_F[\varepsilon] = e_H) + (1 - p_F)\mathcal{L}(0, \mathbb{E}_F[\varepsilon] = 0) = p_F e_H^2, \tag{13}$$

where we have utilized (12). It is straightforward to see that for $a - \sigma(1 + a) > 0$ any deviation $m < 0$ is strictly inferior to $-m > 0$. Thus we consider only deviations with $m > 0$ in the following. A deviation $m > 0$ always results in expectations $\mathbb{E}_F[\varepsilon] = e_H$. Consequently, the most profitable of these deviations is m_H^T. In line with (12), losses for this deviation are

$$\mathcal{L}\left(m_H^T, \mathbb{E}_F[\varepsilon] = e_H\right) = p_F^* e_H^2. \tag{14}$$

There is no profitable deviation for H if $p_F^* e_H^2 \geq p_F e_H^2$ (compare (13) and (14)) or, equivalently, $p_F \leq p_F^*$. Due to the symmetry of the firms' optimization problem, this also implies that no profitable deviation exists for $-H$ in this case.

Deviations for L and -L We show next that no profitable deviation exists for L. Again it suffices to examine deviations with $m > 0$, as any deviation $m < 0$ is strictly inferior to $-m > 0$ for $a - \sigma(1+a) > 0$. Type L's equilibrium losses are $p_F \mathcal{L}(0, e_L) + (1 - p_F)\mathcal{L}(0, 0) = p_F e_L^2$, while a deviation m with $m > 0$ entails losses $p_F \mathcal{L}(m, e_L) + (1 - p_F)\mathcal{L}(m, e_H)$. The most profitable of all deviations is $m_{\mathcal{E}}^T$ with $\mathcal{E} := p_F e_L + (1 - p_F) e_H$.[17]

This deviation will not be attractive if the equilibrium losses are smaller than the losses incurred by choosing $m_{\mathcal{E}}^T$

$$p_F e_L^2 < p_F \mathcal{L}(m_{\mathcal{E}}^T, e_L) + (1 - p_F)\mathcal{L}(m_{\mathcal{E}}^T, e_H),$$

which, utilizing (12), can be expressed as

$$p_F e_L^2 < p_F \left[p_F^* e_L^2 + (1 - p_F^*)(\mathcal{E} - e_L)^2\right] + (1 - p_F)\left[p_F^* e_H^2 + (1 - p_F^*)(\mathcal{E} - e_H)^2\right].$$

This inequality always holds, because $p_F e_L^2 < p_F p_F^* e_L^2 + (1 - p_F) p_F^* e_H^2$ for $p_F \leq p_F^*$. Hence there is no profitable deviation for L. The demonstration that a profitable deviation does not exist for $-L$ is completely analogous.

To sum up, no type $\tau \in \mathcal{T}$ has a profitable deviation, and the equilibrium outlined in the proposition exists.

A.2 Uniqueness

Next, we show that no other equilibria exist. Again we focus on the case $a - \sigma(1+a) > 0$ and omit the case with $a - \sigma(1+a) < 0$, which is completely analogous. In line with

[17] This fact can be easily checked by solving the respective first-order condition for m.

our additional assumptions about $f(m)$, L and H will choose weakly positive values of m in any equilibrium. Moreover, $-L$ and $-H$ will choose weakly negative values of m. In addition, the monotonicity of $f(m)$ implies the monotonicity of the central bank's decisions as a function of its private estimate of the shock, which can be formally stated as $0 = m_0^O \leq m_L^O \leq m_H^O$ (and $m_{-H}^O \leq m_{-L}^O \leq m_0^O = 0$).

These considerations entail that three constellations are possible with regard to L and H, in addition to the constellation we have already considered ($m = 0$ for both). First, L may choose 0, and H may choose a strictly positive m ($0 = m_L^O < m_H^O$). Second, L may choose a strictly positive value for m that is strictly lower than H's choice ($0 < m_L^O < m_H^O$). Third, both types may pool ($0 < m_L^O = m_H^O$). Therefore uniqueness can be established by ruling out all three additional constellations.

First we demonstrate that no equilibrium with $0 = m_L^O < m_H^O$ exists. As H separates itself from the other types, $f(m_H^O) = e_H$ must hold. We have already demonstrated that H could profitably deviate to $m = 0$ if $m_H^O = m_H^T$ and $p_F < p_F^*$. If m_H^O were different from m_H^T, deviating would be even more profitable for $p_F < p_F^*$ because m_H^T is the value of m that minimizes H's losses under the restriction that $f(m) = e_H$.

Second, we consider $0 < m_L^O < m_H^O$. In a fully separating equilibrium, $f(m_L^O) = e_L$ and $f(m_H^O) = e_H$ must hold. Again, H could profitably deviate if $p_F < p_F^*$. Thus no separating equilibrium exists.

Third, it remains to be shown that semi-separating equilibria with $0 < m_L^O = m_H^O$ can be ruled out. For such an equilibrium $f(m_L^O) = (\rho_L e_L + \rho_H e_H)/(\rho_L + \rho_H) = 2(\rho_L e_L + \rho_H e_H) =: \hat{\mathcal{E}} > e_L$ must hold. We introduce e_L^O as the solution to $m_{e_L^O}^T = m_L^O$ (compare (9)), i.e. $e_L^O := \left[(\sigma^2 + a(1-\sigma)^2) / (a - \sigma(1+a))\right] m_L^O$. In line with (12), type L's losses in the semi-separating equilibrium would amount to

$$\begin{aligned}
\mathcal{L}_{L,1} &:= p_F \mathcal{L}(m_L^O, e_L) + (1 - p_F)\mathcal{L}(m_L^O, \hat{\mathcal{E}}) \\
&= p_F \left(p_F^* e_L^2 + (1 - p_F^*)(e_L^O - e_L)^2\right) + (1 - p_F) \left(p_F^* \hat{\mathcal{E}}^2 + (1 - p_F^*)(e_L^O - \hat{\mathcal{E}})^2\right)
\end{aligned}$$

For the deviation to $m = 0$, type L's losses would be

$$p_F \mathcal{L}(m = 0, \mathbb{E}_F[\varepsilon] = e_L) = p_F e_L^2.$$

The equilibrium does not exist if $\mathcal{L}_{L,1} > p_F e_L^2$ holds $\forall e_L^O$. To show this condition, we note that $p_F e_L + (1 - p_F)\hat{\mathcal{E}}$ is the value of e_L^O that minimizes $\mathcal{L}_{L,1}$. Evaluating $\mathcal{L}_{L,1}$ at $e_L^O = p_F e_L + (1 - p_F)\hat{\mathcal{E}}$ yields

$$\begin{aligned}\mathcal{L}_{L,1} &= p_F^*(p_F e_L^2 + (1 - p_F)\hat{\mathcal{E}}^2) \\ &\quad + (1 - p_F^*)(p_F(p_F e_L + (1 - p_F)\hat{\mathcal{E}} - e_L)^2 + (1 - p_F)(p_F e_L + (1 - p_F)\hat{\mathcal{E}} - \hat{\mathcal{E}})^2) \\ &= p_F^*(p_F e_L^2 + (1 - p_F)\hat{\mathcal{E}}^2) \\ &\quad + (1 - p_F^*)(p_F(1 - p_F)^2(e_L - \hat{\mathcal{E}})^2 + (1 - p_F)p_F^2(e_L - \hat{\mathcal{E}})^2) \\ &= p_F^*(p_F e_L^2 + (1 - p_F)\mathcal{E}^2) + (1 - p_F^*)p_F(1 - p_F)(e_L - \hat{\mathcal{E}})^2.\end{aligned}$$

The difference between $\mathcal{L}_{L,1}$, evaluated at $e_L^O = p_F e_L + (1 - p_F)\hat{\mathcal{E}}$, and the losses incurred by deviating to $m = 0$, $p_F e_L^2$, can be readily computed as

$$\mathcal{L}_{L,1} - p_F e_L^2 = p_F^*(1 - p_F)\hat{\mathcal{E}}^2 + (1 - p_F^*)\left[p_F(1 - p_F)(\hat{\mathcal{E}} - e_L)^2 - p_F e_L^2\right].$$

This expression is always positive for $\hat{\mathcal{E}} > e_L$ and $p_F \leq p_F^*$. Consequently, for type L a profitable deviation always exists, and semi-separating equilibria with $0 < m_L^O = m_H^O$ can be ruled out.

Hence we have established existence and uniqueness of the equilibrium outlined in the proposition. □

B Proof of Proposition 3

As a first step, we specify beliefs and, in particular, out-of-equilibrium beliefs. We have to distinguish between $a - \sigma(1 + a) > 0$ and $a - \sigma(1 + a) < 0$. For $a - \sigma(1 + a) > 0$ beliefs are

$$f(m) = \begin{cases} -e_H & \text{for } m < m_{-L}^T \\ -e_L & \text{for } m_{-L}^T \leq m < 0 \\ 0 & \text{for } m = 0 \\ +e_L & \text{for } 0 < m \leq m_L^T \\ +e_H & \text{for } m > m_L^T \end{cases} \tag{15}$$

and for $a - \sigma(1+a) < 0$ they are

$$f(m) = \begin{cases} e_H & \text{for } m < m_L^T \\ e_L & \text{for } m_L^T \leq m < 0 \\ 0 & \text{for } m = 0 \\ -e_L & \text{for } 0 < m \leq m_{-L}^T \\ -e_H & \text{for } m > m_{-L}^T. \end{cases}$$

For the remainder of the proof we assume $a - \sigma(1+a) > 0$. Adapting the proof to $a - \sigma(1+a) < 0$ is straightforward. In the following, we have to prove that no profitable deviation exists for all types $\tau \in \mathcal{T}$.

Deviations for 0 It is easy to show that 0 cannot profitably deviate from $m = 0$. Even if $f(m) = 0 \ \forall m$ held, $m = 0$ would be preferable over and against all $m \neq 0$. For all m with $f(m) \neq 0$, type 0's losses would be even higher than in the case where $f(m) = 0$ would hold. Thus $m = 0$ represents the optimal choice for the beliefs defined in (15).

Deviations for -H and H It suffices to consider only possible deviations of H, as the analysis of type $-H$'s deviations is completely analogous. It is important to note that a deviation with $m < 0$ always leads to higher losses than $-m > 0$. Therefore we focus on deviations with $m > 0$.

According to (15), all deviations m with $m > m_L^T$ entail $f(m) = e_H$. As m_H^T is H's optimal choice, conditional on $f(m) = e_H$, these deviations are not profitable.

A deviation to 0 implies $f(0) = 0$. We note that $\hat{p}_F > p_F^*$. Thus $p_F > \hat{p}_F$ implies $p_F > p_F^*$. According to the proof of Proposition 2, type H therefore prefers m_H^T with $f(m_H^T) = e_H$ to 0 with $f(m_H^T) = 0$. Hence $m = 0$ never represents a profitable deviation.

Finally, we have to check whether deviating to a value of m with $0 < m \leq m_L^T$ might yield lower losses to H. For such a deviation, $f(m) = e_L$ according to (15). The most profitable of these deviations is m_L^T.[18] Thus we need to compare type H's losses for m_H^T

[18]Conditional on $f(m) = e_L$, $m = p_F m_H^T + (1 - p_F) m_L^T$ would minimize losses. Because this value

with its losses for m_L^T. If a central bank of type H chooses m_H^T, losses can be computed using (12):

$$\mathcal{L}(m_H^T, e_H) = p_F^* e_H^2$$

By contrast, if H chooses m_L^T, its losses will amount to

$$\begin{aligned} \mathcal{L}_{H,3} &:= p_F \mathcal{L}(m_L^T, e_H) + (1 - p_F)\mathcal{L}(m_L^T, e_L) \\ &= p_F^*(p_F e_H^2 + (1 - p_F)e_L^2) + p_F(1 - p_F^*)(e_H - e_L)^2. \end{aligned} \tag{16}$$

As a consequence, there is no profitable deviation if $\mathcal{L}_{H,3} \geq \mathcal{L}(m_H^T, e_H)$ or

$$p_F^*(p_F e_H^2 + (1 - p_F)e_L^2) + p_F(1 - p_F^*)(e_H - e_L)^2 \geq p_F^* e_H^2.$$

This inequality can be reformulated as

$$p_F \left[p_F^*(e_H^2 - e_L^2) + (1 - p_F^*)(e_H - e_L)^2\right] \geq p_F^* e_H^2 - p_F^* e_L^2.$$

Rearranging terms and applying $e_H^2 - e_L^2 = (e_H - e_L)(e_H + e_L)$ yields

$$p_F \geq \frac{p_F^* e_H^2 - p_F^* e_L^2}{p_F^*(e_H^2 - e_L^2) + (1 - p_F^*)(e_H - e_L)^2} = \frac{e_H + e_L}{e_H + (2p_F^* - 1)e_L} p_F^* = \hat{p}_F.$$

Hence, if $p_F \geq \hat{p}_F$, there is no profitable deviation for H. Otherwise a profitable deviation exists.

Deviations for -L and L Again we focus on deviations of L with $m \geq 0$. According to the proof of Proposition 2, deviating to $m = 0$ is not profitable if $p_F > p_F^*$, which holds because of $p_F > \hat{p}_F$ and $\hat{p}_F > p_F^*$. Choosing a value of m from the interval $]0; m_L^T[$ is never profitable, as this would entail $f(m) = e_L$ and m_L^T is the value of m that minimizes type L's losses contingent on $f(m) = e_L$. It remains to be shown that L cannot lower its losses by choosing $m > m_L^T$. Such a choice implies $f(m) = e_H$. The deviation with $m > m_L^T$ that yields the lowest losses can be easily computed as $m_{\mathcal{E}}^T$ with $\mathcal{E} = p_F e_L + (1 - p_F)e_H$. Following (12), this deviation implies losses

$$\begin{aligned} &\mathcal{L}_{L,3} := p_F \mathcal{L}(m_{\mathcal{E}}^T, e_L) + (1 - p_F)\mathcal{L}(m_{\mathcal{E}}^T, e_H) \\ &\quad = p_F \left[p_F^* e_L^2 + (1 - p_F^*)(e_L - \mathcal{E})^2\right] + (1 - p_F)\left[p_F^* e_H^2 + (1 - p_F^*)(e_H - \mathcal{E})^2\right] \\ &\quad = p_F \left[p_F^* e_L^2 + (1 - p_F^*)(1 - p_F)^2(e_H - e_L)^2\right] + (1 - p_F)\left[p_F^* e_H^2 + (1 - p_F^*)p_F^2(e_H - e_L)^2\right] \\ &\quad = p_F^*(p_F e_L^2 + (1 - p_F)e_H^2) + p_F(1 - p_F)(1 - p_F^*)(e_H - e_L)^2. \end{aligned} \tag{17}$$

of m is strictly larger than m_L^T, $\text{argmin}_{m \in]0; m_L^T]}\{p_F \mathcal{L}(m, e_H) + (1 - p_F)\mathcal{L}(m, e_L)\} = m_L^T$.

In equilibrium, L's losses are

$$\mathcal{L}(m_L^T, e_L) = p_F^* e_L^2.$$

Thus from L's perspective deviating is not desirable if $\mathcal{L}_{L,3} \geq \mathcal{L}(m_L^T, e_L)$, which is equivalent to

$$p_F^*(1-p_F)(e_H^2 - e_L^2) + p_F(1-p_F)(1-p_F^*)(e_H - e_L)^2 > 0.$$

As this inequality always holds, all deviations lead to higher losses for L over and against the equilibrium losses. Consequently, we have demonstrated that the proposed equilibrium exists for $p_F \geq \hat{p}_F$. For $p_F < \hat{p}_F$, the equilibrium does not exist because $-H$ and H can profitably deviate to m_L^T in this case. □

C Proof of Proposition 4

As in previous proofs, we focus on $a - \sigma(1+a) > 0$. The analysis of the case with $a - \sigma(1+a) < 0$, which is completely analogous, is omitted.

Deviations for 0 Again, no profitable deviation exists for type 0, because $m = 0$ is the money growth this type of central bank would also choose under transparency and any deviation results in beliefs $f(m) = e_H$, which involves even higher costs than the same deviation would entail for $f(m) = 0$.

Deviations for H and -H We focus on type H and note that it is straightforward to extend the analysis to $-H$. As m_H^T is the optimal choice, given that $f(m) = e_H$, no deviation with $m > 0$ can ever be profitable. Moreover, we do not have to examine deviations with $m < 0$ because $-m > 0$ would always be strictly more desirable in these cases. Consequently, we only have to check the deviation $m = 0$, which involves $f(m) = 0$. In the proof of Proposition 2 we have already demonstrated that choosing $m = 0$ is not profitable compared to $m = m_H^T$ if $f(0) = 0$, $f(m_H^T) = e_H$, and $p_F \geq p_F^*$.

Deviations for L and -L Again we omit the analysis of type $-L$'s deviations. For type L, deviations with $m < 0$ are never desirable. The most profitable of all deviations with $m > 0$ is $p_F m_L^T + (1-p_F) m_H^T$. This choice leads to losses $\mathcal{L}_{L,3}$ (see (17)). In equilibrium, L's losses are

$$p_F \mathcal{L}(0, e_L) + (1-p_F)\mathcal{L}(0,0) = p_F e_L^2.$$

Deviating is not attractive if $\mathcal{L}_{L,3} > p_F e_L^2$ or

$$p_F^*(p_F e_L^2 + (1-p_F)e_H^2) + p_F(1-p_F)(1-p_F^*)(e_H - e_L)^2 > p_F e_L^2. \tag{18}$$

This inequality holds for $p_F = p_F^*$ and is violated for $p_F = 1$. Moreover, we note that the difference between the left-hand side and the right-hand side of the inequality is quadratic in p_F. Consequently, there is a unique value of p_F, denoted by $\widetilde{p}_F$ ($p_F^* < \widetilde{p}_F < 1$) such that (18) holds with equality. Hence the proposed equilibrium exists for $p_F^* \leq p_F \leq \widetilde{p}_F$. □

D Proof of Proposition 5

As a first step, we specify beliefs and, in particular, out-of-equilibrium beliefs. We have to distinguish between $a - \sigma(1+a) > 0$ and $a - \sigma(1+a) < 0$. For $a - \sigma(1+a) > 0$ beliefs are

$$f(m) = \begin{cases} -e_H & \text{for } m < \phi m_{-L}^T \\ -e_L & \text{for } \phi m_{-L}^T \leq m < 0 \\ 0 & \text{for } m = 0 \\ +e_L & \text{for } 0 < m \leq \phi m_L^T \\ +e_H & \text{for } m > \phi m_L^T \end{cases} \tag{19}$$

and for $a - \sigma(1+a) < 0$ they are

$$f(m) = \begin{cases} e_H & \text{for } m < \phi m_L^T \\ e_L & \text{for } \phi m_L^T \leq m < 0 \\ 0 & \text{for } m = 0 \\ -e_L & \text{for } 0 < m \leq \phi m_{-L}^T \\ -e_H & \text{for } m > \phi m_{-L}^T. \end{cases}$$

For the remainder of the proof we assume $a - \sigma(1+a) > 0$. Adapting the proof to $a - \sigma(1+a) < 0$ is straightforward. In the following, we examine the conditions under which no profitable deviation exists for all types.

Deviations of -L and L It again suffices to study only possible deviations of L, as the extension of the analysis to type $-L$'s deviations is straightforward. In equilibrium, L chooses ϕm_L^T, and the attentive firms' beliefs amount to e_L, irrespective of whether they have received direct information or have inferred the size of the shock from the central bank's policy. Applying (12) yields the losses in equilibrium:

$$\mathcal{L}_{L,4} := p_F^* e_L^2 + (1 - p_F^*)(1 - \phi)^2 e_L^2 \tag{20}$$

We note that no deviations with $0 < m < \phi m_L^T$ are ever profitable, as $\phi m_L^T < m_L^T$ holds and m_L^T is the most profitable choice if $f(m) = e_L$. Given $f(m) = e_H$, $m_{\mathcal{E}}^T$ with $\mathcal{E} = p_F e_L + (1 - p_F)e_H$ is the most profitable option. Hence it is sufficient to check only two candidate deviations, namely 0 and $m_{\mathcal{E}}^T$.

According to (12), deviation 0 involves losses

$$\mathcal{L}_{L,5} := p_F \mathcal{L}(0, e_L) + (1 - p_F)\mathcal{L}(0, 0) = p_F e_L^2$$

and, in line with (17), deviation $m_{\mathcal{E}}^T$ results in

$$\mathcal{L}_{L,3} = p_F^*(p_F e_L^2 + (1 - p_F)e_H^2) + p_F(1 - p_F)(1 - p_F^*)(e_H - e_L)^2. \tag{21}$$

Interestingly, $\mathcal{L}_{L,5} > \mathcal{L}_{L,3}$ follows from the fact that (18) is violated for $p_F > \widetilde{p}_F$. Thus $m_{\mathcal{E}}^T$ represents the most profitable of all deviations, and condition $\mathcal{L}_{L,3} \geq \mathcal{L}_{L,4}$ alone guarantees that L cannot profitably deviate.

As a next step, we examine the range of ϕ for which $\mathcal{L}_{L,3} \geq \mathcal{L}_{L,4}$ holds. For $\phi = 0$, this inequality is violated due to ($\mathcal{L}_{L,4} = e_L^2 > \mathcal{L}_{L,5} > \mathcal{L}_{L,3}$), and for $\phi = 1$ it holds strictly. Consequently, there is a unique value of $\phi \in]0; 1[$ that satisfies $\mathcal{L}_{L,3} = \mathcal{L}_{L,4}$. We use $\underline{\phi}$ to denote this value. If and only if $\phi \geq \underline{\phi}$, no profitable deviation exists for L.

Deviations of -H and H In equilibrium, H's losses are $\mathcal{L}(m_H^T, e_H) = p_F^* e_H^2$. For H, $m = 0$ cannot represent a profitable deviation, because $p_F \geq p_F^*$ (compare the proof of Proposition 2). No deviation to $m > \phi e_L$ can be profitable. These deviations imply $f(m) = e_H$, and m_H^T is the optimal choice in this case. Given $f(m) = e_L$, $p_F m_H^T + (1 - p_F)m_L^T$ would be optimal, which is larger than ϕm_L^T. Consequently, ϕm_L^T

is the most profitable of all deviations $m \in]0; \phi m_L^T]$. If H deviates to ϕm_L^T, losses will be

$$\mathcal{L}_{H,4} := p_F \left[p_F^* e_H^2 + (1-p_F^*)(e_H - \phi e_L)^2 \right] + (1-p_F) \left[p_F^* e_L^2 + (1-p_F^*)(e_L - \phi e_L)^2 \right]. \tag{22}$$

No profitable deviation for H exists if $\mathcal{L}_{H,4} \geq \mathcal{L}(m_H^T, e_H)$. We note that $\mathcal{L}_{H,4} = p_F e_H^2 + (1-p_F)e_L^2 > p_F^* e_H^2 = \mathcal{L}(m_H^T, e_H)$ for $\phi = 0$ and $\mathcal{L}_{H,4} < \mathcal{L}(m_H^T, e_H)$ for $\phi = 1$ (which follows from $p_F < \hat{p}_F$). As a result, there is a value of $\phi \in]0; 1[$ with $\mathcal{L}_{H,4} = \mathcal{L}(m_H^T, e_H)$, which will be denoted by $\overline{\phi}$. For every $\phi \leq \overline{\phi}$ type H cannot profitably deviate.

Does $\underline{\phi} < \overline{\phi}$ hold? Finally, we have to show $\underline{\phi} < \overline{\phi}$ for all $p_F \in]\widetilde{p}_F, \hat{p}_F[$. For this purpose, we demonstrate $\mathcal{L}_{H,4} > \mathcal{L}(m_H^T, e_H)$ at $\phi = \underline{\phi}$.

Recall that $\underline{\phi}$ is defined by $\mathcal{L}_{L,3} = \mathcal{L}_{L,4}$. Applying (20) yields

$$\mathcal{L}_{L,3} = p_F^* e_L^2 + (1-p_F^*)(1-\underline{\phi})^2 e_L^2. \tag{23}$$

With the help of (16) and (23), $\mathcal{L}_{H,4}$ (in (22)) can be written as

$$\begin{aligned}
\mathcal{L}_{H,4} = p_F &\left[p_F^* e_H^2 + (1-p_F^*) \left(e_H - e_L + (1-\underline{\phi}) e_L \right)^2 \right] \\
&+ (1-p_F) \left[p_F^* e_L^2 + (1-p_F^*)(1-\underline{\phi})^2 e_L^2 \right] \\
= p_F &\left[p_F^* e_H^2 + (1-p_F^*) \left\{ (e_H - e_L)^2 + 2(e_H - e_L)(1-\underline{\phi}) e_L + (1-\underline{\phi})^2 e_L^2 \right\} \right] \\
&+ (1-p_F) \left[p_F^* e_L^2 + (1-p_F^*)(1-\underline{\phi})^2 e_L^2 \right] \\
= \mathcal{L}_{H,3} &+ 2p_F(1-p_F^*)(e_H - e_L)(1-\underline{\phi}) e_L + (1-p_F^*)(1-\underline{\phi})^2 e_L^2 \\
= \mathcal{L}_{H,3} &+ 2p_F(1-p_F^*)(e_H - e_L)(1-\underline{\phi}) e_L + \mathcal{L}_{L,3} - p_F^* e_L^2.
\end{aligned}$$

Hence $\mathcal{L}_{H,4} > \mathcal{L}(m_H^T, e_H) = p_F^* e_H^2$ holds at $\phi = \underline{\phi}$ if $\mathcal{L}_{H,3} + \mathcal{L}_{L,3} > p_F^*(e_H^2 + e_L^2)$. Using (16) and (21), this condition can be easily verified. □

E Proof of Proposition 6

E.1 Case $\mathbf{p_F < p_F^*}$

For $p_F < p_F^*$, the statement of the proposition is a direct consequence of the proof of Proposition 2. There we have shown that each type of bank $\tau \in \mathcal{T} \setminus \{0\}$ prefers 0 with

$f(0) = 0$ to m_τ^T with $f(m_\tau^T)$, provided that $p_F < p_F^*$. Thus each of these central-bank types has lower losses under opacity than under transparency. Moreover, type 0's losses are unaffected by the transparency regime. Consequently, expected social losses are lower under opacity for $p_F < p_F^*$.

E.2 Case $\mathbf{p_F > p_F^*}$

The case with $p_F > p_F^*$ is more intricate, because the equilibria under opacity are not unique in general. We proceed by showing that any potential equilibrium under opacity yields higher losses compared to the transparency solution. While it is unclear for which parameter constellations these potential equilibria exist (if they exist at all), we prove that, if they existed, they would definitely lead to higher social losses over and against the equilibrium under transparency.

First we demonstrate that semi-separating equilibria with $m_L^O = 0$ and $m_H^O > 0$ can never be superior to the transparency solution. In the proof of Proposition 2, we have shown that type L's losses are lower for m_L^T than for 0 if $f(m_L^T) = e_L$, $f(0) = 0$, and $p_F > p_F^*$. Type H's losses can never be lower in a semi-separating equilibrium with $m_L^O = 0$ compared to transparency, as the money growth chosen under transparency minimizes losses, given that $f(m) = e_H$. Obviously, losses for type 0 are identical under transparency and opacity. As losses are weakly higher under opacity for all types and strictly higher for some, expected social losses are strictly lower under transparency.

Second, we note that the fully separating equilibrium where all central banks choose the same money growth rates as under transparency is the fully separating equilibrium with the lowest social losses. Consequently, social losses under transparency are weakly lower over and against opacity under any fully separating equilibrium. For $p_F^* < p_F < \hat{p}_F$, they are strictly lower under transparency irrespective of which equilibrium is chosen under opacity. This is a consequence of Proposition 3, which states that equilibria where all types make the same choices as under transparency do not exist under opacity for $p_F < \hat{p}_F$.

Third, it remains to be shown that semi-separating equilibria with $0 < m_L^O = m_H^O$ would yield higher losses compared to the transparency solution. For such a semi-separating equilibrium $f(m_L^O) = f(m_H^O) = (\rho_L e_L + \rho_H e_H)/(\rho_L + \rho_H) = 2(\rho_L e_L + \rho_H e_H) = \hat{\mathcal{E}}$, where we have utilized $\rho_H + \rho_L = 1/2$. In equilibrium, type L's losses would amount to

$$\mathcal{L}_{L,1} = p_F \mathcal{L}(m_L^O, e_L) + (1 - p_F)\mathcal{L}(m_L^O, \hat{\mathcal{E}})$$

and type H's losses would be

$$\mathcal{L}_{H,1} := p_F \mathcal{L}(m_L^O, e_H) + (1 - p_F)\mathcal{L}(m_L^O, \hat{\mathcal{E}}).$$

Under transparency, type L's losses are

$$\mathcal{L}_{L,2} := \mathcal{L}(m_L^T, e_L) = p_F^* e_L^2$$

and type H's losses are given by

$$\mathcal{L}_{H,2} := \mathcal{L}(m_H^T, e_H) = p_F^* e_H^2.$$

A semi-separating equilibrium where L and H pool would yield higher expected social costs than the transparency solution, if $\rho_L(\mathcal{L}_{L,1} - \mathcal{L}_{L,2}) + \rho_H(\mathcal{L}_{H,1} - \mathcal{L}_{H,2}) > 0 \; \forall m$. It is straightforward to verify that the value of m_L^O that minimizes the left-hand side of this inequality is $m_{\hat{\mathcal{E}}}^T$. Evaluating the left-hand side of the above inequality for this value, we obtain

$$\begin{aligned}
& \rho_L(\mathcal{L}_{L,1} - \mathcal{L}_{L,2}) + \rho_H(\mathcal{L}_{H,1} - \mathcal{L}_{H,2}) \\
= \; & \rho_L \left[p_F \left(p_F^* e_L^2 + (1 - p_F^*) \left(e_L - \hat{\mathcal{E}} \right)^2 \right) + (1 - p_F) p_F^* \hat{\mathcal{E}}^2 - p_F^* e_L^2 \right] \\
& \quad + \rho_H \left[p_F \left(p_F^* e_H^2 + (1 - p_F^*) \left(e_H - \hat{\mathcal{E}} \right)^2 \right) + (1 - p_F) p_F^* \hat{\mathcal{E}}^2 - p_F^* e_H^2 \right] \\
= \; & \frac{1}{2}(1 - p_F) p_F^* \hat{\mathcal{E}}^2 + \rho_L \left[p_F (1 - p_F^*) \left(e_L - \hat{\mathcal{E}} \right)^2 - (1 - p_F) p_F^* e_L^2 \right] \\
& \quad + \rho_H \left[p_F (1 - p_F^*) \left(e_H - \hat{\mathcal{E}} \right)^2 - (1 - p_F) p_F^* e_H^2 \right] \\
= \; & \frac{1}{2}(1 - p_F) p_F^* \hat{\mathcal{E}}^2 + \rho_L \left[p_F (1 - p_F^*) \left(e_L^2 - 2 e_L \hat{\mathcal{E}} + \hat{\mathcal{E}}^2 \right) - (1 - p_F) p_F^* e_L^2 \right] \\
& \quad + \rho_H \left[p_F (1 - p_F^*) \left(e_H^2 - 2 e_H \hat{\mathcal{E}} + \hat{\mathcal{E}}^2 \right) - (1 - p_F) p_F^* e_H^2 \right]
\end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{2}(1-p_F)p_F^*\hat{\mathcal{E}}^2 + (p_F - p_F^*)\left[\rho_L e_L^2 + \rho_H e_H^2\right] \\
&\qquad + p_F(1-p_F^*)\left[\frac{1}{2}\hat{\mathcal{E}}^2 - 2(\rho_L e_L + \rho_H e_H)\hat{\mathcal{E}}\right] \\
&= \frac{1}{2}(1-p_F)p_F^*\hat{\mathcal{E}}^2 + (p_F - p_F^*)\left[\rho_L e_L^2 + \rho_H e_H^2\right] - \frac{1}{2}p_F(1-p_F^*)\hat{\mathcal{E}}^2 \\
&= (p_F - p_F^*)\left[\rho_L e_L^2 + \rho_H e_H^2 - 2\rho_H^2 e_H^2 + 4\rho_L\rho_H e_L e_H - 2\rho_L^2 e_L^2\right] \\
&= (p_F - p_F^*)\left[\rho_L(1-2\rho_L)e_L^2 + \rho_H(1-2\rho_H)e_H^2 + 4\rho_L\rho_H e_L e_H\right] \\
&= 2(p_F - p_F^*)\rho_H\rho_L(e_H + e_L)^2,
\end{aligned}
$$

where we have applied (12), $\hat{\mathcal{E}} = 2(\rho_L e_L + \rho_H e_H)$, and $\rho_L + \rho_H = \frac{1}{2}$. $2(p_F - p_F^*)\rho_H\rho_L(e_H + e_L)^2$ is positive if $p_F > p_F^*$. Consequently, transparency yields strictly lower losses than any semi-separating equilibrium that might exist.

To sum up, transparency is strictly superior to opacity for $p_F^* < p_F < \hat{p}_F$. For $p_F \geq \hat{p}_F$, transparency and opacity will be equivalent with regard to welfare if the equilibrium outlined in Proposition 3 is chosen. If another equilibrium is chosen, transparency will be strictly more desirable than opacity from the aggregate welfare perspective. □

F Proof of Proposition 7

In this variant of our model, the attentive firms' expectations about the shock will be identical to the shock's realization if these firms receive information about the shock either directly or indirectly from the central bank. If attentive firms do not receive such information, their expectations about the shock will be influenced by the central bank's monetary policy in general because it is possible that the central bank withholds its information.

We prove that, if $p_F \leq p_F^*$, an equilibrium exists in which the central bank never withholds its information. We start with a description of the attentive firms' beliefs about the central bank's type. Later we will see that these beliefs are actually consistent with the central bank's equilibrium strategy. Suppose that attentive firms expect $\mathbb{E}_F[\varepsilon] = 0$ if $m = 0$. For positive money growth rates, they expect $\mathbb{E}_F[\varepsilon] = e_H$ if $a - \sigma(1+a) > 0$ and $\mathbb{E}_F[\varepsilon] = -e_H$ if $a - \sigma(1+a) < 0$. For negative money growth rates, they expect $\mathbb{E}_F[\varepsilon] = -e_H$ if $a - \sigma(1+a) > 0$ and $\mathbb{E}_F[\varepsilon] = +e_H$ if $a - \sigma(1+a) < 0$. Loosely speaking, if the central bank appears to respond to a shock despite claiming that it does not know about a shock, firms will surmise that the central bank is hiding information about a large shock whose sign is consistent with the direction of the central bank's response.

In the candidate equilibrium in which the central bank releases its information irrespective of its type, it is obvious that central banks of all types $\tau \in \mathcal{T}$ will choose policy m_τ^T. It remains to be shown that no profitable deviation for types $\mathcal{T} \setminus \{0\}$ exists. According to our analysis of the transparency scenario, no profitable deviation to an $m \neq m_\tau^T$ exists when the central bank announces its private information. Thus we have to examine whether a deviation exists for which the central bank withholds its information. It is straightforward to verify that no such deviation with $m \neq 0$ can be profitable. This leaves the combination of $m = 0$ and withholding information about $\varepsilon = e_\tau$ as the only possible profitable deviation. In this case, the central bank's losses will be $\mathcal{L}(0, e_\tau) = p_F e_\tau^2$ as opposed to $\mathcal{L}(m_\tau^T, e_\tau) = p_F^* e_\tau^2$, which are the losses in the candidate equilibrium. As a consequence, the deviation will not be beneficial to the central bank if $p_F \geq p_F^*$. The attentive firms' beliefs are consistent with the central bank's equilibrium strategy. Hence the equilibrium exists if $p_F \geq p_F^*$. □

References

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

George-Marios Angeletos and Alessandro Pavan. Efficient Use of Information and Social Value of Information. *Econometrica*, 75(4):1103–1142, July 2007.

Romain Baeriswyl and Camille Cornand. Monetary Policy and its Informative Value. paper presented at the conference on "Cycles, Contagion and Crises" at LSE, London, on 28-29 June 2007, October 2007.

Romain Baeriswyl and Camille Cornand. The Signaling Role of Policy Action. *Journal of Monetary Economics*, 2010. forthcoming.

Laurence Ball, N. Gregory Mankiw, and Ricardo Reis. Monetary Policy for Inattentive Economies. *Journal of Monetary Economics*, 52(4):703–725, May 2005.

Olivier Blanchard and Jordi Galí. Real Wage Rigidities and the New Keynesian Model. *Journal of Money, Credit, and Banking*, 39(s1):35–65, February 2007.

Olivier Blanchard and Jordi Galí. Labor Markets and Monetary Policy: A New-Keynesian Model with Unemployment. mimeo, March 2008.

Alan S. Blinder, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen. Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence. *Journal of Economic Literature*, 46(4):910–945, December 2008.

Willem H. Buiter. Alice in Euroland. *Journal of Common Market Studies*, 73(2): 181–209, 1999.

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

Antonello D'Agostino and Karl Whelan. Federal Reserve Information During the Great Moderation. *Journal of the European Economic Association*, 6(2-3):609–620, 2008.

Jon Faust and Lars E. O. Svensson. Transparency and Credibility: Monetary Policy with Unobservable Goals. *International Economic Review*, 42(2):369–397, 2001.

Petra M. Geraats. Central Bank Transparency. *Economic Journal*, 112(483):532–565, 2002.

Hans Gersbach and Volker Hahn. Information Content of Wages and Monetary Policy. *Journal of Money, Credit, and Banking*, 39(1):133–149, February 2007.

Hans Gersbach and Volker Hahn. Voting Transparency in a Monetary Union. *Journal of Money, Credit, and Banking*, 41(5):809–830, August 2009.

Marvin Goodfriend. Monetary Mystique: Secrecy and Central Banking. *Journal of Monetary Economics*, 17(1):63–92, January 1986.

Volker Hahn. Transparency in Monetary Policy: A Survey. *ifo Studien*, 48(3):429–455, 2002.

Volker Hahn. Transparency in Monetary Policy, Signaling, and Heterogeneous Information. mimeo, October 2009.

Christian Hellwig. Heterogeneous Information and the Benefits of Transparency. mimeo, February 2005.

Henrik Jensen. Optimal Degrees of Transparency in Monetary Policymaking: The Case of Imperfect Information about the Cost-Push Shock, November 2000. mimeo.

Henrik Jensen. Optimal Degrees of Transparency in Monetary Policymaking. *Scandinavian Journal of Economics*, 104(3):399–422, 2002.

Karen Lewis. Why Doesn't Society Minimize Central Bank Secrecy? *Economic Inquiry*, 39:99–112, 1991.

N. Gregory Mankiw and Ricardo Reis. Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics*, 117(4):1295–1328, 2002.

Stephen Morris and Hyun Song Shin. Social Value of Public Information. *American Economic Review*, 92(5):1521–1534, 2002.

Christina D. Romer and David H. Romer. Federal Reserve Information and the Behavior of Interest Rates. *American Economic Review*, 90(3):429–457, June 2000.

Anne Sibert. Monetary Policy with Uncertain Central Bank Preferences. *European Economic Review*, 46:1093–1109, 2002.

Anne Sibert. Monetary Policy Committees: Individual and Collective Reputation. *Review of Economic Studies*, 70(3):649–665, 2003.

Anne Sibert. Is Transparency about Central Bank Plans Desirable? *Journal of the European Economic Association*, 7(4):831857, June 2009.

Michael Spence. Job Market Signaling. *The Quarterly Journal of Economics*, 87(3): 355–374, August 1973.

Lars E. O. Svensson. Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro-Transparency, Not Con. *American Economic Review*, 96(1): 448–452, March 2006.

John Vickers. Signalling in a Model of Monetary Policy with Incomplete Information. *Oxford Economic Papers*, 38(3):443–455, November 1986.

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

Michael Woodford. Inflation Stabilization and Welfare. *Contributions to Macroeconomics*, 2(1):1009–1009, 2002.