Oxford Economic Papers 64 (2012), 593–615 **593**
doi:10.1093/oep/gpr056

Should central banks remain silent about their private information on cost-push shocks?

By Volker Hahn

Center of Economic Research, ETH Zurich, ZUE D13, 8092 Zurich, Switzerland; e-mail: vhahn@ethz.ch

We propose a signalling 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, the central bank should remain silent about its 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 signalling incentives under opacity. An *ex ante* transparency requirement can improve welfare even if central banks have the possibility to withhold information discretionarily. Moreover, our model may provide a rationale for the recent trend towards more transparency in monetary policy.

JEL classifications: D82, D83, E52, E58.

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]

---

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

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

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. The difference between the abovementioned finding in Angeletos and Pavan (2007) and ours can be explained in the following way. In both papers, 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. An additional effect arises in the present paper. If the central bank does not publish its private information, the firms will attempt to infer this information by observing monetary policy. As a consequence, the central bank has to take into account the information that it signals by its policy. This may lead to signalling costs and, in particular, to distorted responses to moderate shock realizations. These distortions result from the fact that, upon observing moderate shocks, the central bank has to make sure that firms do not expect it to react to large 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, the central bank may have the discretionary power to withhold information. If this was actually beneficial to the central bank, a transparency requirement would 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 (e.g., Buiter, 1999), 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 shocks 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. Some robustness issues and proofs are relegated to online Appendices 4-9.

2. Related literature

Our paper contributes to two strands of the literature. First, it is part of the literature on signalling games, which goes back to Spence (1973). In monetary economics, signalling games have been studied by Vickers (1986), Sibert (2002, 2003), Gersbach and Hahn (2007, 2009), and Sibert (2009). 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 expectations. 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 signalling 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 literature considers the economic effects of central-bank communication as well as the consequences that the publication of private central-bank

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

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 (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 in the present paper is affected by firms' expectations about the shock rather than the actual shock realization.

Our paper is also related to Walsh (2007) and Baeriswyl and Cornand (2007, 2010). 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.

3. Model

We consider a one-shot signalling 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

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

[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.

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

The central bank observes the shock with probability p_{CB} $(0 \leqslant p_{CB} \leqslant 1)$. With the complementary probability, the central bank obtains no private information. Consequently, there exist five types of central bank. 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

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

[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.44-5).

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

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

where $a \geqslant 0$ 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 probability p_F $(0 \leqslant p_F \leqslant 1)$. The event of them observing the shock is 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:

(i) Nature draws the shock ε'.
(ii) The central bank becomes informed about ε' with probability p_{CB}.
(iii) Under opacity, the central bank's information is kept private. Under transparency, it is published.
(iv) The central bank chooses its instrument m.
(v) 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.
(vi) All firms choose the prices of their outputs.

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 $\mathbb{E}_F$:

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

[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 Section 9, we discuss the assumption that the central bank's instrument can be observed.

Solving for p yields

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

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:

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

where

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

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 $\mathbb{E}_F[\varepsilon] = e_\tau$ $(\tau \in \mathcal{T} \setminus \{0\})$. Then m^T_τ represents the optimal choice of type $\tau \in \mathcal{T} \setminus \{0\}$ 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 type's perspective, namely $-e_H$, $-e_L$, 0, e_L, or 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

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

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 $m_0^T := 0$.

We summarize our observations in the following proposition:

Proposition 1 Under transparency, a unique equilibrium exists. Each type of central bank $\tau\in\mathcal{T}$ chooses m_τ^T.

In this equilibrium, the 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 signalling 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 $-e_H\leqslant f(m)\leqslant e_H\ \forall m$ 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 $\tau\in\mathcal{T}$ are a weakly monotonic function of the central bank's estimate of the shock. Formally, this implies $m^O_{-H}\leqslant m^O_{-L}\leqslant m^O_0=0\leqslant m^O_{+L}\leqslant m^O_{+H}$ for $a-\sigma(1+a)>0$ and $m^O_{-H}\geqslant m^O_{-L}\geqslant m^O_0=0\geqslant m^O_{+L}\geqslant m^O_{+H}$ for $a-\sigma(1+a)<0$, where we use 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 $0\leqslant p^*_F<1$.[13]

We are now in a position to describe the equilibria under opacity.

Proposition 2 For $p_F<p^*_F$, a unique[14] equilibrium exists. In this equilibrium, all types of central bank $\tau\in\mathcal{T}$ 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 $f(m)=+e_H\ \forall m<0$.

---

[12]The third consequence of our assumptions can be explained as follows. Suppose, for example, $a-\sigma(1+a)>0$. Then $f(m)\geqslant 0\ \forall m>0$ 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 $|m|$ 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 $m_H\geqslant m_L$ (otherwise the firms' beliefs would be incorrect). In a similar vein, $m_{-H}\leqslant m_{-L}$ follows from the monotonicity of $f(m)$.

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

[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.

The proof is given in Appendix 1 (existence) and Appendix 5 (uniqueness). 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 demonstrate that an equilibrium satisfying our additional assumptions always exists by characterizing equilibria for different intervals of p_F. 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 $p_F \geqslant \hat{p}_F$, there is an equilibrium under opacity in which all types of central bank $\tau \in \mathcal{T}$ choose the money growth rates they would find optimal under transparency (m_τ^T).

For the proof, see Appendix 2. 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 that $\hat{p}_F > p_F^*$. Consequently, we have to show that perfect Bayesian Nash equilibria satisfying our additional assumptions on $f(m)$ exist for the interval $p_F^* < p_F < \hat{p}_F$. Intuitively, separating equilibria with the same choices as under transparency do not exist, as there would be strong incentives for a central bank 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 behaviour is given in the following proposition, proved in Appendix 6:

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 $p_F \in [p_F^*, \widetilde{p}_F]$. A central bank of type $\tau \in \{-L, 0, +L\}$ chooses $m = 0$. A central bank of type $\tau \in \{-H, +H\}$ chooses 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 ($p_F \geqslant \hat{p}_F$), separating equilibria exist where all central-bank types display the same behaviour under opacity as under transparency. Moreover, for sufficiently small values of p_F ($p_F \leqslant p_F^*$), pooling equilibria exist. For somewhat larger p_F ($p_F^* \leqslant p_F \leqslant \tilde{p}_F$), semi-separating equilibria occur in which central banks of types $-L$ and L mimic the behaviour of zero. We note that $\tilde{p}_F < \hat{p}_F$ cannot be ruled out, as can be readily verified. Thus it remains to describe possible equilibria for the interval $p_F \in]\tilde{p}_F; \hat{p}_F[$. This gap is filled by the following proposition, which is proved in Appendix 7:

Proposition 5 Suppose $\tilde{p}_F < \hat{p}_F$. For all $p_F \in]\tilde{p}_F; \hat{p}_F[$, 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 $m^T_0 = 0$.

These equilibria are particularly interesting as they represent fully separating equilibria where the behaviour 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 behaviour 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 3, 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 $p_F \geqslant \hat{p}_F$, 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 $f(m)$.

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

[15]For $p_F \geqslant \hat{p}_F$, 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.

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 signalling 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 signalling 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 has observed a shock that it has not observed. Moreover, a central bank that does not observe a shock cannot prove that it does not possess information and thus cannot certify its type.[16] We will call transparency feasible if a perfect Bayesian Nash equilibrium of this game of partially certifiable information exists in which the central bank never withholds its private information.

Such an equilibrium of the game where the central bank can decide on transparency after observing the shock will be a separating equilibrium (if it exists) because the attentive firms can exactly identify the central bank's type. This is reminiscent of the equilibria of the game under opacity that are described in Propositions 3 and 5, which also represent separating equilibria. However, in the latter cases the attentive firms learn the central bank's type from observing monetary policy rather than from the information that the central bank decides to publish after observing it, which is the case studied in this section. In Appendix 8, we show

[16]For an overview of games with certifiable information see Bolton and Dewatripont (2005, ch.5). Games of partially certifiable information are studied by Shin (1994), among others.

Proposition 7 Transparency is feasible if $p_F \geqslant p_F^*$.

A comparison with Proposition 6 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.

Why does the social desirability of transparency entail that it is feasible? Committing the central bank to transparency is socially beneficial when opacity would involve signalling costs. In this case, the central bank also benefits from avoiding these signalling costs by releasing its private information discretionarily. For example, a central bank having observed L prefers to prove its type, thus preventing the attentive firms from mistaking it for a central bank of type H.

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 number of instruments and shocks, the additional restrictions on equilibrium we have introduced under opacity, the observability of the instrument, the structure of firms' information, imprecise signals, and the possibility of reputation-building in an infinitely repeated version of our game.

9.1 Specification of shocks

In this paper we have focused on four different shock realizations. This number is sufficiently high to identify the important signalling 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 signalling 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.[17]

9.2 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 examine demand shocks instead, but transparency regarding these shocks would never be socially harmful. Under opacity, separating equilibria

[17] A detailed analysis is available upon request.

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.

9.3 Two shocks at the same time

As a next step, consider our basic model and assume that there are demand shocks in addition to cost-push shocks. Accordingly, aggregate demand is $y = m - p + \mu$ for some arbitrarily distributed demand shock μ. For simplicity of exposition, we make the assumption that the central bank observes the demand shock with certainty but the firms receive no information about its size.

Then our analysis can be used to address the question whether the central bank should publish information on cost-push shocks, given that it is already transparent about demand shocks. This can be seen by defining $\tilde{m} := m + \mu$, which can be interpreted as the difference between the instrument actually chosen by the central bank (m) and the choice of instrument ($-\mu$) that would be necessary to stabilize only the demand shock. Importantly, all propositions would continue to hold if we substituted $\tilde{m}$ for m. For example, if the probability of firms observing cost-push shocks was low ($p_F < p_F^*$), an equilibrium would exist in which $\tilde{m} = 0$ would be chosen by the central bank, irrespective of its information about the cost-push shock. Intuitively, the central bank would perfectly stabilize demand shocks ($\tilde{m} = 0$ entails $m = -\mu$) but would be passive with respect to cost-push shocks.[18]

9.4 Restrictions on equilibrium

In Section 6, where we analyse 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.

9.5 Observability of instrument

As has been explained before, the fact that the central bank cannot separate shock stabilization from signalling information when adjusting its instrument under opacity is crucial for our findings. This link relies on the assumption that the

[18]The observation that the central bank can perfectly stabilize demand shocks lends plausibility to our previous assumption that it chooses to be transparent with respect to these shocks.

attentive firms can observe the central bank's instrument. It is therefore warranted to discuss the plausibility of this assumption in more detail.

In practice, central banks' ability to keep their instruments secret is limited. The Federal Reserve did not publish its policy directive before 1994. However, target changes could be inferred from movements in the federal funds rate rather quickly. Cook and Hahn (1989, p.332) relate that, from 1974 to 1979, the Fed's control of the federal funds rate was so firm that the public could detect most target changes on the day they occurred. The financial press reported these changes on the following day. Interestingly, the signalling effect of monetary policy actions was identified by Romer and Romer (2000) for a period during which the policy directive was not published immediately. To sum up, if the central bank cannot ensure full secrecy of its instrument, the only possible approach to severing the link between shock stabilization and signalling information is transparency about shocks, which is examined in this paper.

9.6 Structure of firms' information

It is instructive to discuss our assumptions on the firms' possible information sets. In our model, there are two different channels of information that firms may utilize under opacity, namely the central bank's decisions and own research about the size of aggregate shocks. While information about monetary policy can be easily obtained from the financial press, forming an own estimate of aggregate shocks is arguably more difficult. It thus seems natural to assume that all firms that observe own information about shocks also use readily available information on monetary policy.

How would our results change if we modified this assumption? The distortions under opacity arise from the incentives of types H and $-H$ to signal moderate shock realizations and thus to mimic the behaviour of types L and $-L$. These incentives depend on the possibility that firms' estimates of the shock may be influenced by their observations of the instrument. This possibility arises as long as there is a positive probability of some firms observing the instrument but not independent information about the shock. Allowing for the additional possibility of some other firms observing own information but not the central bank's decisions would not eliminate the incentives to signal information, which are crucial for our results regarding welfare.

9.7 Modelling imprecise information

In our paper, imperfect information about the shock is modelled by the assumption that the central bank and the firms observe the shock only with a positive probability. An alternative approach to modelling imperfect information about shocks is to assume that a signal is observed with certainty but that this signal is imprecise (see the literature on central-bank transparency surveyed in Section 2). Our assumption has been made for analytical convenience, as it facilitates the computation of expectations about the shock and the expectations about other agents' observations. For example, an attentive firm that observes information about the shock does not have to infer the central bank's information from the instrument under opacity because the central bank's signal contains no information over and against the firm's own information.

Our results could be directly extended to the case where the central bank's and attentive firms' signals are noisy, provided that the noise terms were perfectly correlated. Introducing imperfect correlation between the central bank's signal and the signal observed by firms would severely complicate the analysis but would plausibly lead to even higher signalling costs under opacity. In this case, also the firms with an own signal would try to infer the central bank's signal from its policy, as the central bank's signal would contain additional information about the shock. Hence imperfectly correlated signals would tighten the link between shock stabilization and information signalling that, as we have argued before, is crucial for our result that transparency may be beneficial.[19]

9.8 Repetition of the game

If we repeated our game infinitely many times, the equilibria of the one-shot game could also be used to characterize all Markov-perfect equilibria of this new game. In these equilibria, the central bank's strategy prescribes an optimal choice in each period, given the central bank's information and the behaviour of firms. Hence monetary policy is credible in the sense that firms know that the central bank will never deviate from its strategy. All of our results about the relative merits of transparency and opacity would continue to hold in this case. Under opacity, there would be additional equilibria in which firms play strategies that are not Markov perfect. In some of these equilibria, reputation-building may alleviate the distortions arising under opacity for $p_F > p_F^*$. Plausibly, this might be achieved for sufficiently high discount factors if firms could coordinate on a trigger strategy. This trigger strategy would punish an opaque central bank in the future if it was detected to choose its instrument in a way that would not be optimal under transparency. However, it is not clear that such an equilibrium would be played and so the same outcome as under transparency could be attained under opacity.[20]

[19]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 the central bank is always informed about cost-push shocks.

[20]Moreover, equilibria in which the distortions under opacity are eliminated by reputation-building would not exist if the shocks were drawn from a distribution with an unbounded support. Intuitively, there would always be a positive probability of a shock realization that would be so large such that the central bank would prefer the short-term gains from deviating from the policy that would be optimal under transparency over the long-term costs incurred by losing reputation. For a related argument, see footnote 6 in Loisel (2008), who analyses reputation-building in a New Keynesian framework.

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 central banks 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 signalling incentives of different types of central bank 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.

Supplementary material

Supplementary material (Appendices 4-9) is available at the OUP website.

Acknowledgements

I would like to thank Johannes G. Becker, Hans Gersbach, Oliver Grimm, Stephan Imhof, Maik Schneider, participants of the EEA annual meeting in Glasgow, 2010, two anonymous referees, and the editors for many valuable comments and suggestions. A previous version of this paper was circulated under the title 'Why the Publication of Socially Harmful Information May Be Socially Desirable'.

References

Adam, K. (2007) Optimal monetary policy with imperfect common knowledge, *Journal of Monetary Economics*, **54**, 267–301.

Angeletos, G.-M. and Pavan, A. (2007) Efficient use of information and social value of information, *Econometrica*, **75**, 1103–42.

Baeriswyl, R. and Cornand, C. (2007) Monetary policy and its informative value, Paper presented at the conference on 'Cycles, Contagion and Crises' at LSE, London, 28-29 June 2007.

Baeriswyl, R. and Cornand, C. (2010) The signaling role of policy action, *Journal of Monetary Economics*, **57**, 682–95.

Ball, L., Mankiw, N.G., and Reis, R. (2005) Monetary policy for inattentive economies, *Journal of Monetary Economics*, **52**, 703–25.

Blinder, A.S., Ehrmann, M., Fratzscher, M., De Haan, J., and Jansen, D.-J. (2008) Central bank communication and monetary policy: a survey of theory and evidence, *Journal of Economic Literature*, **46**, 910–45.

Bolton, P. and Dewatripont, M. (2005) *Contract Theory*, Cambridge, MA, MIT Press.

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

Cook, T. and Hahn, T. (1989) The effect of changes in the Federal Funds rate target on market interest rates in the 1970s, *Journal of Monetary Economics*, **24**, 331–51.

Cukierman, A. and Meltzer, A.H. (1986) A theory of ambiguity, credibility, and inflation under discretion and asymmetric information, *Econometrica*, **54**, 1099–128.

D'Agostino, A. and Whelan, K. (2008) Federal Reserve information during the Great Moderation, *Journal of the European Economic Association*, **6**, 609–20.

Faust, J. and Svensson, L.E.O. (2001) Transparency and credibility: Monetary policy with unobservable goals, *International Economic Review*, **42**, 369–97.

Geraats, P.M. (2002) Central bank transparency, *Economic Journal*, **112**, 532–65.

Gersbach, H. and Hahn, V. (2007) Information content of wages and monetary policy, *Journal of Money, Credit, and Banking*, **39**, 133–49.

Gersbach, H. and Hahn, V. (2009) Voting transparency in a monetary union, *Journal of Money, Credit, and Banking*, **41**, 809–30.

Hahn, V. (2002) Transparency in monetary policy: a survey, *IFO Studien*, **48**, 429–55.

Hahn, V. (2009) *Transparency in monetary policy, signaling, and heterogeneous information*, mimeo, Center of Economic Research, EHT, Zurich.

Hellwig, C. (2005) *Heterogeneous information and the benefits of transparency*, mimeo, University of California, Los Angeles.

Jensen, H. (2000) *Optimal degrees of transparency in monetary policymaking: the case of imperfect information about the cost-push shock*, mimeo, University of Copenhagen.

Jensen, H. (2002) Optimal degrees of transparency in monetary policymaking, *Scandinavian Journal of Economics*, **104**, 399–422.

Lewis, K. (1991) Why doesn't society minimize central bank secrecy? *Economic Inquiry*, **39**, 99–112.

Loisel, O. (2008) Central bank reputation in a forward-looking model, *Journal of Economic Dynamics and Control*, **32**, 3718–42.

Mankiw, N.G. and Reis, R. (2002) Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve, *Quarterly Journal of Economics*, **117**, 1295–328.

Morris, S. and Shin, H.S. (2002) Social value of public information, *American Economic Review*, **92**, 1521–34.

Romer, C.D. and Romer, D.H. (2000) Federal Reserve information and the behavior of interest rates, *American Economic Review*, **90**, 429–57.

Shin, H.S. (1994) News management and the value of firms, *The RAND Journal of Economics*, **25**, 58–71.

Sibert, A. (2002) Monetary policy with uncertain central bank preferences, *European Economic Review*, **46**, 1093–109.

Sibert, A. (2003) Monetary policy committees: individual and collective reputation, *Review of Economic Studies*, **70**, 649–65.

Sibert, A. (2009) Is transparency about central bank plans desirable? *Journal of the European Economic Association*, **7**, 831–57.

Spence, M. (1973) Job market signaling, *The Quarterly Journal of Economics*, **87**, 355–74.

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**, 448–52.

Vickers, J. (1986) Signalling in a model of monetary policy with incomplete information, *Oxford Economic Papers*, **38**, 443–55.

Walsh, C.E. (2007) Optimal economic transparency, *International Journal of Central Banking*, **3**, 5–36.

Woodford, M. (2002) Inflation stabilization and welfare, *Contributions to Macroeconomics*, **2,** available at http://www.bepress.com/bejm/contributions/vol2/iss1/art1.

Appendix 1: Proof of Proposition 2: existence

In this Appendix, we establish existence. Uniqueness is considered in online Appendix 5. 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)). Suppose that the central bank chooses m_e^T, which is the choice that would be optimal under transparency if the attentive firms believed the shock to be e. Using (8) and (9), we obtain the following expression for the central bank's losses in this case:

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

Deviations for 0 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 examine 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 to a negative value of m cannot be optimal because a positive value of m

of the same size would be superior in this case. 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

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

There is no profitable deviation for H if $p_F^* e_H^2 \geqslant p_F e_H^2$ (compare (13) and (14)) or, equivalently, $p_F \leqslant p_F^*$. 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 positive deviations, as any negative deviation would be dominated by a positive deviation of the same size 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 $\mathcal{E} := p_F e_L + (1 - p_F) e_H$.[21]

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

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

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 $p_F \leqslant p_F^*$. Hence there is no profitable deviation for L. Showing that no profitable deviation exists 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. □

Appendix 2: 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

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

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

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

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

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 0 represents the optimal choice of m, given the beliefs defined in (15).

Deviations for $-H$ and H It suffices to consider possible deviations of H, as the analysis of type $-H$'s deviations is completely analogous. We note that a deviation to a negative value of m always leads to higher losses than a deviation to the respective positive value of the same size. 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 $0<m\leqslant m_L^T$ 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.[22] Thus we need to compare type H's losses for m_H^T with its losses for m_L^T. If a central bank of type H chooses m_H^T, losses can be computed with the help of (12):

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

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

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

[22] Conditional on $f(m)=e_L$, $m=p_Fm_H^T+(1-p_F)m_L^T$ would minimize losses. Because this value of m is strictly larger than m_L^T, $\mathrm{argmin}_{m\in]0;m_L^T]}\{p_F\mathcal{L}(m,e_H)+(1-p_F)\mathcal{L}(m,e_L)\}=m_L^T$.

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

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

This inequality can be reformulated as

$$p_F\big[p_F^*(e_H^2 - e_L^2) + (1-p_F^*)(e_H - e_L)^2\big] \geqslant p_F^* e_H^2 - p_F^* e_L^2.$$

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

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

Hence, if and only if $p_F \geqslant \hat{p}_F$, there is no profitable deviation for H.

Deviations for $-L$ and L Again we focus on deviations of L with $m \geqslant 0$. 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 some $m > m_L^T$. 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

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

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 $\mathcal{L}_{L,3} \geqslant \mathcal{L}(m_L^T, e_L)$, 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 $p_F \geqslant \hat{p}_F$. 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. □

Appendix 3: Proof of Proposition 6

3.1 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 $p_F < p_F^*$.

3.2 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 every potential equilibrium under opacity yields higher losses compared to the transparency solution (the only exception being the equilibrium outlined in Proposition 3, which is equivalent with respect to losses). 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. Checking all potential equilibria is straightforward but tedious. These steps are therefore relegated to online Appendix 9. □