Does Freedom of Information Deter Information Acquisition?

Rossella Argenziano* Helen Weeds†

December 22, 2014

Abstract

We consider the impact of freedom of information laws on incentives for information acquisition. A biased decision-maker chooses whether to acquire costly information to inform his decision regarding a policy action. The decision-maker's private optimum is to act too often, and the private value of information is below the social optimum. After the action choice a monitor chooses whether to investigate, using freedom of information laws to reveal the decision-maker's information. When the cost of acquiring information is low, freedom of information rules which lower the monitor's cost/reward ratio from investigation discipline the decision-maker's action, moving this closer to the social optimum. For intermediate information costs, however, the threat of investigation inhibits information acquisition, and lowering the cost of investigation simply reduces information acquisition with no increase in discipline. In extensions to the main model we consider observable information acquisition and voluntary disclosure. We highlight implications of the analysis for freedom of imformation rules and institutional design.

*University of Essex. Email: rargenz@essex.ac.uk.
†University of Essex and CEPR. Correspondence address: Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK. Email: hfweeds@essex.ac.uk.

1 Introduction

Transparency in decision-making (or 'open government') is often seen as a vital part of a well-functioning democracy. The ability of citizens, journalists and interested parties to scrutinize decisions taken by public bodies is regarded by many commentators as an important mechanism for ensuring that public bodies and officials act in the public interest, rather than following their private desires or pandering to favored groups. Accordingly, many countries have passed Freedom of Information (FOI) laws or open government codes establishing the principle that citizens should be able to access any document held by a public body. A similar mechanism arises in court proceedings: under most well-developed legal systems the parties are required to disclose information to the other side.

Advocates of open government rules extol the advantages of allowing access to information without generally considering the impact of such rules on the incentives for public bodies to gather information. Generally, one might expect public bodies to take steps to acquire information about the circumstances surrounding and consequences of their decisions, in order to assess accurately (or as accurately as possible) the appropriate decision to take. But when information that is gathered might later be subject to a freedom of information request, with any adverse revelation having potentially detrimental or embarrassing consequences for the decision-maker, there might be a 'chilling' effect on information acquisition. This fear may be heightened by the threat of court proceedings or a hostile press. Our paper investigates this possibility.

The paper sets out a model consisting of two players, a biased decision-maker and a monitor. The decision-maker (who may be a politician, regulator or agent) chooses a policy action, deciding whether to act or not, where this choice affects social welfare. The social benefit from the policy choice depends on the state of the world, which is unknown at the start of the game. The decision-maker also receives a private benefit from acting (rather than not acting), creating a wedge between social and private preferences and biasing the decision-maker in favour of acting. Prior to deciding how to act the decision-maker may, at some cost, acquire information revealing the state of the world. The monitor observes the decision-maker's choice of action but does not observe the information acquired (if any), nor even the fact of its acquisition. After observing the decision-maker's action the monitor may, at some cost, open an investigation in which she requests the information acquired by the decision-maker using freedom of information rules. We assume that information is 'hard' in the sense that it cannot be hidden or distorted. If, on the basis of this information, the decision-maker is found to have taken the socially inferior action, the decision-maker incurs a penalty while the monitor receives a reward. In extensions to the baseline model we consider the effect of observable information acquisition, and include the possibility of voluntary information revelation by the decision-maker

The decision-maker's private optimum, without threat of investigation, is affected by bias in two distinct ways. First, the private value of information is below its social value, and hence the private incentive to acquire information is too low. This is the case because bias reduces the impact of the state of the world on the decision-maker's privately optimal action; in the limiting case, the decision-maker's action is determined entirely by his private value from acting and information is of no value to him. Secondly, there is an interval of states of the world over which not acting is the social optimum but an informed decision-maker chooses to act, reflecting his bias in favour of acting.

In the baseline model we find that the threat of investigation alters the decision-maker's behaviour as follows. When the cost of information acquisition is low the decision-maker acquires information, and the intensity of monitoring decreases with the monitor's cost/reward ratio for investigation. In this case the threat of investigation disciplines the action choice of an informed decision-maker somewhat, moving it closer to (but not equal to) the social optimum. This finding provides a basis for the claim that, by lowering investigation costs, freedom of information rules improve decision-making.

When information is more costly, however, we find that the threat of investigation weakens the decision-maker's incentive to acquire information compared with the private optimum. In this interval, moreover, lowering the cost of investigation makes matters worse: there is no increase in monitoring intensity, and hence no improvement in discipline, but the probability of information acquisition falls. Furthermore, in this region making information acquisition observable worsens the situation further still: rather than information being acquired with some positive probability, with observability it is not acquired at all.

Accordingly, our results reveal the possibility of a 'chilling effect' of freedom of information rules on decision-makers' incentives to acquire better information to guide their actions. The impact of freedom of information rules depends crucially on the cost of acquiring information: at low information costs the effect is beneficial, disciplining the actions of decision-makers, but at higher information costs the chilling effect on information acquisition arises. The reduced incentive to acquire information is detrimental to social welfare: uninformed decisions are less likely to achieve the social optimum than biased ones and, accordingly, the quality of decision-making may be reduced overall. Moreover, at higher information costs, measures which strengthen freedom of information by lowering the cost of investigation adversely impact the acquisition of information with no improvement in discipline: far from improving social welfare, such changes actually worsen it.

The possibility of voluntary disclosure tends to accentuate both the disciplining effect on the decision-maker's action and the chilling effect on information acquisition, compared with the baseline model. The tendency for a decision-maker that takes the socially optimal action to disclose his information makes it harder for one that acts contrary to the public interest to hide this fact, as failure to disclose is then more likely to trigger an investigation. But the chilling effect on information acquisition is also heightened: at intermediate levels of information cost the probability that information is acquired is lower.

The paper is related to several strands of literature. One is the literature on the costs and benefits of transparency in the political process, based on principal-agent theory. In a situation of moral hazard with complete contracting, the benchmark result of Holmström (1979) demonstrates that the principal is never harmed by (and generally gains from) observing additional information about the agent's performance. In the absence of complete contracts, however, inefficiencies can arise, as shown by e.g. Maskin and Tirole (2004). A number of papers consider the impact of transparency using models of career concerns. In a theoretical analysis, Prat (2005) distinguishes between two types of information that the principal may have about the agent: information about the agent's action and information about the consequences of the action. He finds that, while transparency on consequences is beneficial, transparency on actions may cause the agent to disregard useful information, choosing actions that are optimal based on public information despite private signals that would suggest otherwise. Accordingly Prat recommends that actions should not be revealed before their consequences are observed, and claims that freedom of information rules in many jurisdictions respect this principle by allowing short-term secrecy while the decision process is ongoing. In an adverse selection model where the agent can gather private information before the principal offers the contract, Hoppe (2013) finds that the principal may be better off when information gathering is a hidden action.

The extension of our model in which the decision maker has the option of voluntarily disclosing his information is related to the literature on persuasion: see Milgrom (1981) and Milgrom and Roberts (1986) and, more recently, Mathis (2008), Che and Kartik (2009), Rayo and Segal (2010), Kamenica and Gentzkow (2011, 2014), Felgenhauer and Loerke (2013, 2014), Felgenhauer and Schulte (2014) and Hagenbach, Koessler and Perez-Richet (2014); see also Milgrom (2008) for a survey. In common with papers on persuasion, and in contrast to 'cheap talk', in our analysis information is taken to be 'hard' and cannot be distorted. However, our paper differs from the persuasion literature in two respects. First, information has a private value to the sender (here, the decision-maker), who himself takes an action, distinct from its value in communication with the receiver (here, the monitor). Secondly, the receiver faces a cost of responding to the sender (investigation is costly), as a result of which the unraveling argument that typically guarantees full disclosure of information in persuasion games does not apply to our model.

The characteristics of freedom of information legislation internationally are described by Frankel (2001), while its effects on UK government are discussed by Hazell and Glover (2011), Hazell, Bourke and Worthy (2012) and Worthy and Hazell (2013). In a theoretical analysis, Levy (2007) considers the impact of transparency on decision-making in committees when members are motivated by career concerns.

The impact of transparency on the decision-making process is studied in a number of empirical papers, drawing on the natural experiment provided by the release of transcripts of the Federal Open Market Committee (FOMC) in 1993: see Meade and Stasavage (2008) and Hansen, McMahon and Prat (2014). Both of these papers find that transparency reduces dissent from the Chairman's policy proposal, increasing the conformity of opinions, while Hansen et al. also find evidence indicating greater information acquisition between meetings by inexperienced members of the committee.

The paper is structured as follows. Section 2 sets out the baseline model. After deriving the decision-maker's private optimum, we introduce a monitor who may use freedom of information rules to investigate the decision-maker's action and solve the strategic game between the two. Section 3 considers two variants of the baseline model: first, where information acquisition by the decision-maker is observable to the monitor, and secondly, including the possibility of voluntary disclosure by the decision-maker. Section 4 discusses the results, drawing implications for freedom of information legislation and institutional design. Longer proofs are presented in the Appendix.

2 The model

2.1 Benchmark: Information acquisition with no monitor

We start by setting out a model of information acquisition by a biased decision-maker in the absence of monitoring. Then, in the next subsection, we introduce a monitor who may open an investigation, using 'freedom of information' rules to compel the decision-maker to reveal his information. The possibility of voluntary information disclosure by the decision-maker is introduced in Section 3.2.

There is a decision-maker, D, who has a decision to make: either to act ($a = 1$, the 'high' action) or not to act ($a = 0$, the 'low' action). The social payoff from the action choice depends on the state of the world θ; the realisation of θ is initially unknown to all parties but its distribution is known to be $\theta \sim U[0, 1]$. Social welfare from action a is given by

$$W = -|a - \theta| \tag{1}$$

where $a \in \{0, 1\}$. Depending on the realisation of θ the socially optimal decision rule is

$$a^* (\theta) = \begin{cases} 1 & \text{if } \theta > \frac{1}{2} \\ 0 & \text{if } \theta < \frac{1}{2} \\ \{0, 1\} & \text{if } \theta = \frac{1}{2}. \end{cases} \tag{2}$$

If the realisation of θ is unknown then the optimal action is assessed according to the prior distribution; as in the final row of (2) society is then indifferent between $a = 1$ and $a = 0$.

Although the decision-maker takes social welfare into account, he also receives a private benefit b from acting ($a = 1$), which may be described as *bias*. The decision-maker's payoff is given by

$$U = ab - |a - \theta| \tag{3}$$

where $b \in (0, 1]$. Accordingly, the decision-maker's relative benefit from $a = 1$ as compared with $a = 0$ is

$$b - 1 + 2\theta. \tag{4}$$

Before making his action choice, the decision-maker chooses whether or not to acquire better information about the state of the world. It is assumed that, by paying a cost k, the decision-maker may learn the realisation of θ.

To summarise, the timeline is as follows:

1. Nature chooses the state of the world θ; this is unobservable.

2. D chooses whether or not to pay k to learn the realisation of θ.

3. D chooses action $a \in \{0, 1\}$; payoffs are realised.

First we consider the optimal action choice of an uninformed decision-maker. In the absence of information acquisition the decision-maker maximises his expected utility given the prior that $E(\theta) = \frac{1}{2}$. The uninformed decision-maker's optimal choice is set out in Lemma 1.

Lemma 1 *In the absence of a monitor, if the decision-maker does not acquire information then he chooses action* $a = 1$*.*

Proof. The result follows from the observation that, given the prior $E(\theta) = \frac{1}{2}$, D's expected relative benefit from acting compared with not acting, given by (4), is $b > 0$, thus D chooses $a = 1$. ■

Given his action choice, the uninformed decision-maker's expected utility is

$$EU_0 = \int_0^1 [b - (1 - \theta)]\, d\theta = b - \frac{1}{2},$$

while expected welfare is

$$EW_0 = \int_0^1 -(1 - \theta)\, d\theta = -\frac{1}{2}.$$

Next we consider the optimal action choice of an informed decision-maker. If the decision-maker knows θ he then makes his decision conditionally, choosing $a = 1$ if and only if expression (4) is weakly positive (we assume that in the case of indifference the decision-maker chooses to act). The decision rule of an informed decision-maker is set out in Lemma 2.

Lemma 2 *In the absence of a monitor, if the decision-maker acquires information then he acts according to*

$$a = \begin{cases} 1 & \text{if } \theta \geq \theta_0 \\ 0 & \text{if } \theta < \theta_0, \end{cases} \tag{5}$$

where $\theta_0 \equiv \frac{1-b}{2}$*.*

Proof. D's expected relative benefit from acting compared with not acting, given by (4), is weakly positive for $\theta \geq \frac{1-b}{2}$ and strictly negative for $\theta < \frac{1-b}{2}$. ■

θ_0 is the privately optimal 'type threshold' describing the action choice of an informed decision-maker in the absence of monitoring. Notice that θ_0 is less than one-half (the socially optimal decision rule): due to the decision-maker's bias in favour of acting there is an interval $\left[\theta_0, \frac{1}{2}\right)$ over which he chooses $a = 1$ despite this being contrary to the social optimum. Also notice that there is no corresponding interval in which the decision-maker chooses $a = 0$ when $a = 1$ would be the social optimum: $a = a^*$ for $\theta \geq \frac{1}{2}$. The informed decision-maker's decision rule in the absence of monitoring is illustrated in Figure 1.

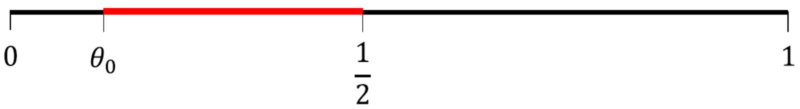

FIGURE 1: D's privately optimal decision rule

The expected utility of an informed decision-maker, ignoring the cost of acquiring information, is given by

$$EU_1 = \int_0^{\theta_0} -\theta d\theta + \int_{\theta_0}^{1} (b - (1 - \theta))\, d\theta = \frac{1}{4}\left(b^2 + 2b - 1\right). \tag{6}$$

Hence the private value of information to the decision-maker in the absence of monitoring is given by

$$EU_1 - EU_0 = \frac{(1-b)^2}{4} \equiv k_0. \tag{7}$$

Accordingly, in the absence of monitoring, information acquisition is optimal for the decision-maker if and only if $k \leq k_0$, as illustrated in Figure 2.

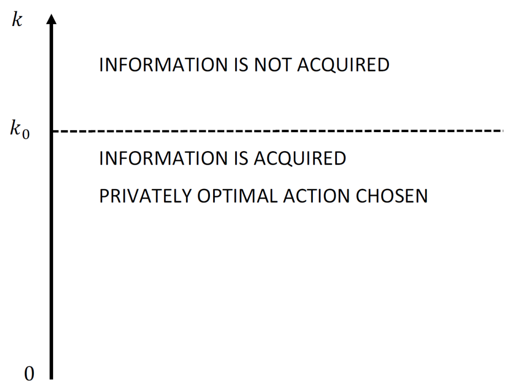

FIGURE 2: D's privately optimal investment threshold

Finally, we consider the social value of information, taking account of the decision-maker's choice of action. Ignoring information acquisition cost k, if the decision-maker acquires information and follows threshold θ_0 then expected welfare is given by

$$EW_1 = \int_0^{\theta_0} -\theta d\theta + \int_{\theta_0}^{1} -(1-\theta)\, d\theta = -\frac{(1+b^2)}{4}.$$

Hence the social value of information is

$$EW_1 - EW_0 = \frac{1-b^2}{4} \equiv k^*.$$

Notice that, for any strictly positive bias, $k^* > k_0$: the social value of information is larger than its private value. Moreover, greater bias reduces the private value of information: with greater bias, information about θ makes less difference to the decision-maker's choice of action. In the limit where $b = 1$, the decision-maker acts regardless of the realisation of θ and information has zero private value.

These findings are summarised in the following proposition.

Proposition 1 (Information acquisition by biased decision-maker)
In the absence of monitoring, greater bias reduces the decision-maker's incentive to acquire information. Compared with the social optimum, the decision-maker invests too little in acquiring information.

2.2 Information acquisition with monitoring

We now introduce a monitor, M. After the decision-maker has chosen his action the monitor may, at a cost $c > 0$, open an investigation in which she invokes freedom of information rules to compel the decision-maker to reveal his information.

The following game is played by the decision-maker and monitor. In this baseline model *information acquisition is unobserved by the monitor.*

1. Nature chooses the state of the world θ; this is unobservable.

2. D chooses to acquire information with probability $\gamma \in [0,1]$, paying k to learn the realisation of θ; this choice is unobserved by M.

3. D chooses an observable action $a \in \{0,1\}$.

4. M chooses whether or not to pay c to open an investigation; payoffs are realised.

The equilibrium concept used throughout the paper is perfect Bayesian equilibrium (PBE). We denote by $\underline{\pi}$ the probability that that the monitor opens an investigation after observing action $a = 0$ and by $\overline{\pi}$ the probability that the monitor opens an investigation after observing $a = 1$.

The investigation process takes place as follows. If the decision-maker has indeed acquired information, then the monitor observes θ. If the decision-maker has not acquired information then the monitor has no other means of learning θ and remains uninformed. If investigation reveals that the decision-maker acted contrary to the socially optimal decision rule (2), i.e. that $a \neq a^*$, then a punishment $p > 0$ is imposed on the decision-maker and the monitor receives a reward $r > 0$. We assume that p and r are fixed amounts, which may differ in magnitude (i.e. punishment and reward are not necessarily a transfer between the players).[1] We assume that $p < b$, i.e. punishment only partially deters biased action.

Normalising to zero the monitor's payoff if she does not open an investigation,[2] her payoff from investigating is given by

$$U^M = r\left|a - a^*(\theta)\right| - c.$$

For notational convenience, we denote the cost/reward ratio $c/r \equiv C$.

The following lemma describes the role of investigation when the decision-maker is uninformed.

[1] An alternative approach would be for the monitor to reverse the decision-maker's action, in which case p and r would be the players' respective utilities from this change.

[2] Depending on the identity of the monitor, she may or may not take account of social welfare resulting from the action itself: a judge, for example, may take full account of social welfare while a journalist might care only about her private cost and reward from monitoring. However, in either case the decision to investigate does not affect the action that has already been taken and this element of the monitor's payoff is the same whether she decides to investigate or not. Therefore equating the payoff from no investigation to zero is just a normalisation.

Lemma 3 *If the decision-maker is uninformed then he chooses $a = 1$. If the monitor believes with probability one that the decision-maker is uninformed then she does not investigate.*

Proof. When D is uninformed, investigation does not reveal the realisation of θ and M bases her assessment on the prior, $E(\theta) = \frac{1}{2}$. Then, from (2), $a^* = \{0, 1\}$: both actions are (weakly) socially optimal and, regardless of his action choice, D receives no punishment. With no threat of punishment, D acts according to his private optimum, shown by Lemma 1 to be $a = 1$. If M believes with probability one that D is uninformed (i.e. she believes that $\gamma = 0$) then, since investigation would incur a cost and have a null expected reward, she chooses not to investigate. ■

The following lemma highlights how the presence of a monitor can only reduce, and never increase, the decision-maker's value of information as compared to the benchmark case with no monitoring.

Lemma 4 *With the possibility of monitoring, the private value of information is weakly less than in the benchmark case.*

Proof. Lemma 3 shows that if D remains uninformed then he is never punished, regardless of which action he chooses, hence his expected payoff in the monitoring game is the same as in the benchmark case. If D acquires information then a punishment may be imposed in some cases, hence his payoff in the continuation game with monitoring is weakly smaller than in the benchmark case. ■

A corollary of Lemma 4 is that, in the game with monitoring, information is never acquired for costs greater than k_0.

2.3 Equilibrium in the monitoring game

The full characterisation of equilibrium in the monitoring game is given in Proposition 2, with a formal proof in the Appendix. Here we describe the intuitions underlying equilibrium outcomes, highlighting two important effects of monitoring: the *discipline* *effect* on the decision-maker's action choice and the *chilling effect* on the acquisition of information.

First, consider the action choice of an informed decision-maker facing a monitor whom he expects to investigate after $a = 1$ (respectively, $a = 0$) with probability $\overline{\pi}$ (respectively $\underline{\pi}$). The following lemma describes the optimal action choice in the presence of monitoring.

Lemma 5 *For any $\{\underline{\pi}, \overline{\pi}\}$, an informed decision-maker acts according to*

$$a = \begin{cases} 1 & \text{if } \theta \geq \theta_{\overline{\pi}} \\ 0 & \text{if } \theta < \theta_{\overline{\pi}}, \end{cases} \tag{8}$$

where $\theta_{\overline{\pi}} \equiv \frac{1-b+\overline{\pi}p}{2}$.

Proof. First, suppose that $\theta \geq \frac{1}{2}$. Given M's strategy, the expected payoff from choosing $a = 0$ is $-\theta - \underline{\pi}p$ while the expected benefit from $a = 1$ is $b - (1 - \theta)$. Hence the relative benefit from acting compared with not acting is $b - 1 + 2\theta + \underline{\pi}p$, which is strictly positive. Next, suppose that $\theta < \frac{1}{2}$. Given M's strategy, the expected payoff from choosing $a = 0$ is $-\theta$ while the expected benefit from $a = 1$ is $b - (1 - \theta) - \overline{\pi}p$. Hence the relative benefit from acting compared with not acting is $b - 1 + 2\theta - \overline{\pi}p$, which is strictly positive (negative) for $\theta > (<)\, \theta_{\overline{\pi}}$. Finally, we have assumed that in the case of indifference D chooses $a = 1$. ■

Notice the following features of the decision-maker's action rule in the presence of monitoring. First, it is independent of $\underline{\pi}$, the probability of investigation after $a = 0$. At any θ for which the decision-maker chooses $a = 0$ this is also the social optimum, thus investigation after $a = 0$ never results in punishment and the threat to investigate after $a = 0$ has no effect.

Secondly, $\theta_{\overline{\pi}}(0) = \theta_0$: if the decision-maker expects no investigation after $a = 1$, then he acts according to his privately optimal decision rule (5). But any positive probability of investigation moves the decision-maker's action threshold *closer* to the social optimum: with $\overline{\pi} > 0$, $\theta_{\overline{\pi}}$ is to the right of θ_0, closer to one-half. This is the *discipline effect*: the threat of investigation disciplines the decision-maker's action, reducing the impact of his bias. The discipline effect is strongest when $\overline{\pi} = 1$; the decision-maker's action threshold in this case is denoted $\theta_1 \equiv \frac{1-b+p}{2} \in \left(\theta_0, \frac{1}{2}\right)$. These thresholds are illustrated in Figure 3. Note that, since $p < b$, even when the monitor investigates for sure after $a = 1$ the decision-maker's bias is only partially disciplined and θ_1 lies strictly to the left of one-half.

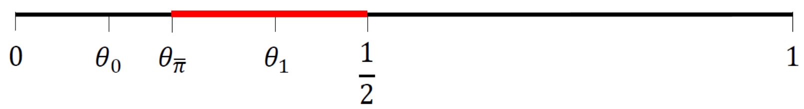

FIGURE 3: D's decision rule with monitoring

Now we turn to the monitor's decision to investigate, which is taken after observing the decision-maker's action. Recall that the monitor cannot observe information acquisition, thus she does not necessarily know whether the decision-maker is informed or not. Her investigation strategy is based on her equilibrium beliefs about information acquisition. The monitor can condition her investigation strategy on the observed action $a \in \{0, 1\}$. After observing $a = 0$, it is optimal for the monitor never to investigate, i.e. she chooses $\underline{\pi} = 0$, as demonstrated by the following lemma.

Lemma 6 *After observing $a = 0$ the monitor chooses not to open an investigation (i.e. $\underline{\pi} = 0$).*

Proof. D is either uninformed or he is informed. If D is uninformed, Lemma 3 shows that it is optimal not to investigate (and, in any case, an uninformed D always chooses $a = 1$). If D is informed, Lemma 5 demonstrates that whenever $a = 0$ is chosen, this is also the socially optimal action (as $\theta_{\overline{\pi}} < \frac{1}{2}$). Hence investigation following $a = 0$ will reveal (in equilibrium) that the action is socially optimal and no reward will be earned. Since investigation following $a = 0$ incurs a cost and has a null expected reward, M chooses not to investigate. ■

Accordingly, M may investigate only after observing $a = 1$. Whether she chooses to do so depends on her expected reward from investigating, relative to its cost, taking account of the probability that the decision-maker is both informed and taking the socially inferior action. As will be shown in Proposition 2 the monitor's investigation strategy depends on the cost-reward ratio C: at low values of C the monitor investigates with high probability, for intermediate values of C the probability of investigation is lower, and for value of C above a critical threshold C_0 she never investigates.

Finally, the decision-maker's decision to acquire information depends on the likelihood of investigation in the subgame when information is acquired (as already noted, when information is not acquired investigation results in no punishment). When C is low, the monitor investigates (after $a = 1$) with high probability and the value of information to the decision-maker is reduced. In consequence, information is acquired only for lower levels of the information cost k. This is the *chilling effect* of freedom of information: by exposing the information-gatherer to the threat of punishment, the incentive to acquire information is inhibited.

Before describing the equilibrium of the game we define the following expressions:

$$\begin{aligned}
C_{\overline{\pi}} &\equiv \frac{b-\overline{\pi}p}{b-\overline{\pi}p+1} \in [C_1, C_0], \text{ where } C_1 \equiv \frac{b-p}{b-p+1}, C_0 \equiv \frac{b}{b+1}; \\
k_1 &\equiv \frac{(1-b)^2 - p(2b-p)}{4} < k_0; \\
k_C &\equiv \frac{1-b}{2} - \frac{1-2C}{4(1-C)^2}; \text{ notice that } k_C(C_1) \equiv k_1 \text{ and } k_C(C_0) \equiv k_0; \\
\overline{\pi}_k &\equiv \frac{1}{p}\left(b - \sqrt{2b+4k-1}\right); \text{ notice that } \overline{\pi}_k(k_0) = 0 \text{ and } \overline{\pi}_k(k_1) = 1; \\
\gamma(k, C) &\equiv \frac{2C}{C + (1-C)\left(\sqrt{2b+4k-1}\right)}.
\end{aligned}$$

Notice that $\gamma(k, 0) = 0$, $\gamma(k_0, C_1) = \frac{2b-2p}{2b-p} \in (0, 1)$ and $\gamma(k_0, C_0) = \gamma(k_1, C_1) = 1$. The following proposition summarises the characterisation of equilibrium in the monitoring game; the proposition is illustrated in Figure 4.

Proposition 2 (Equilibrium in the monitoring game)

- *If $C \in [0, C_1]$ and $k < k_1$, D acquires information and chooses $a = 1$ iff $\theta \geq \theta_1$ and $a = 0$ otherwise. M investigates after observing $a = 1$.*

- *If $C \in [0, \gamma C_1]$ and $k = k_1$, D acquires information with probability $\gamma \in [0, 1]$.*

If uninformed, D chooses $a = 1$; if informed, D chooses $a = 1$ iff $\theta \geq \theta_1$ and $a = 0$ otherwise. M investigates after observing $a = 1$.

- *If $C \in [0, C_0)$ and $k \in [\max(k_1, k_C), k_0]$, D acquires information with probability $\gamma(k, C) \in [0, 1]$. If uninformed, D chooses $a = 1$; if informed, D chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}(k)} \in [\theta_0, \theta_1]$ and $a = 0$ otherwise. After observing $a = 1$ M investigates with probability $\overline{\pi}_k \in [0, 1]$.*

- *If $C \in (C_1, C_0)$ and $k \leq k_C$, D acquires information and chooses $a = 1$ iff $\theta \geq \theta_C \equiv \frac{1-2C}{2(1-C)} \in (\theta_0, \theta_1)$ and $a = 0$ otherwise. After observing $a = 1$ M investigates with probability $\overline{\pi}_C = \frac{1}{p}\left(b - \frac{C}{1-C}\right) \in (0, 1)$.*

- *If $C \in [C_0, 1]$ and $k < k_0$, D acquires information and chooses $a = 1$ iff $\theta \geq \theta_0$ and $a = 0$ otherwise. M never investigates.*

- *If $C \in [\gamma C_0, 1]$ and $k = k_0$, D acquires information with probability $\gamma \in [0, 1]$. If uninformed, D chooses $a = 1$; if informed, D chooses $a = 1$ iff $\theta \geq \theta_0$ and $a = 0$ otherwise. M never investigates.*

- *If $k > k_0$, D never acquires information, and chooses $a = 1$. M never investigates.*

Proof. See Appendix. ■

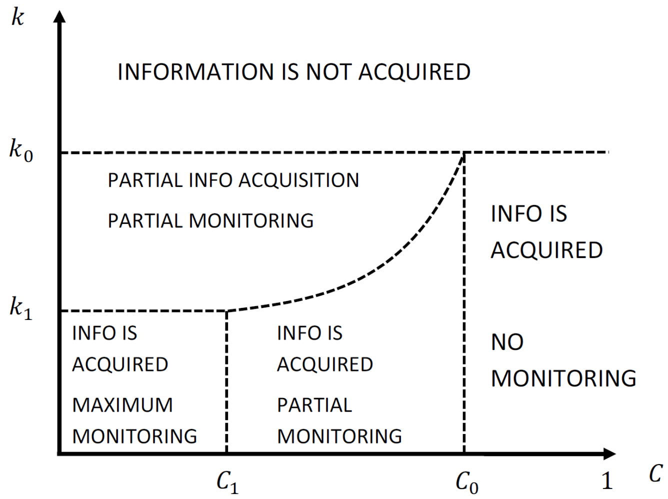

FIGURE 4: Equilibrium in the monitoring game (Proposition 2)

Looking at Figure 4, the impacts of the two cost parameters k and C can be seen as follows. For high information costs $k > k_0$ information is not acquired, the same as in the no-monitoring benchmark. For low information costs $k < k_1$ information is always acquired. This is the case even when the monitor's cost-reward ratio is sufficiently low ($C < C_1$) that investigation always takes place (following $a = 1$): in the bottom left-hand corner there is both full information acquisition and maximum discipline of the decision-maker's action choice ($\theta = \theta_1$). As C increases, monitoring decreases: over the interval (C_1, C_0) the monitor investigates with an (intermediate) probability $\overline{\pi}_C$, which is decreasing in C. The reduction in monitoring intensity increases the value of information to the decision-maker, thus information is acquired at higher information costs: over this interval the threshold for (full) information acquisition is k_C, which is increasing in C. Above C_0 investigation never occurs: the decision-maker acquires information and acts according to his private optimum.

For intermediate values of k between $\max(k_1, k_C)$ and k_0 there is partial information acquisition and partial monitoring. As a mixed strategy equilibrium, the probability with which each player acts ensures indifference on the part of the *other* player. The monitor's probability of investigating after $a = 1$, $\overline{\pi}_k$, is decreasing in k (but independent of C): this ensures the decision-maker's indifference towards information acquisition. The probability that the decision-maker acquires information, $\gamma(k, C)$, is decreasing in k and increasing in C, ensuring that the monitor remains indifferent towards investigation as its cost-reward ratio rises.

Comparative statics differ according to the region into which parameter values fall. If k is sufficiently low that information is always acquired (i.e. for $k < k_C$), increasing the cost-reward ratio C between C_1 and C_0 lowers monitoring intensity and hence reduces the discipline on the decision-maker's action choice. But at intermediate values of k where the mixed equilibrium obtains, increasing C does not affect monitoring intensity but increases information acquisition, thus improving the quality of decision-making. In the same region an increase in information cost k reduces both information acquisition and monitoring intensity, both of which reduce the quality of decision-making. Outside this region increasing k affects neither information acquisition nor monitoring intensity.

3 Extensions

This section examines two variants of the game set out in Section 2. First we make information acquisition by the decision-maker observable to the monitor. Secondly, we vary the original game (with unobservable information acquisition) by giving the decision-maker the option of disclosing his information voluntarily.

3.1 Observable information acquisition

Consider a situation in which the acquisition of information by the decision-maker at stage 2 of the game is observed by the monitor. Proposition 3 summarises the characterisation of equilibrium for this case; this is illustrated in Figure 5.

Proposition 3 (Equilibrium with observable information acquisition)

- *If $C \in [0, C_1]$, D acquires information iff $k \leq k_1$; if $k = k_1$ this may be with probability $\gamma \in [0, 1]$. If uninformed, D chooses $a = 1$; if informed, D chooses $a = 1$ iff $\theta \geq \theta_1$ and $a = 0$ otherwise. M investigates when D is informed and chooses action $a = 1$.*

- *If $C \in (C_1, C_0)$, D acquires information iff $k \leq k_C \in (k_1, k_0)$. If uninformed, D chooses $a = 1$; if informed, D chooses $a = 1$ iff $\theta \geq \theta_C \equiv \frac{1-2C}{2(1-C)} \in (\theta_0, \theta_1)$ and $a = 0$ otherwise. If D is informed and chooses action $a = 1$ M investigates with probability $\overline{\pi}_C = \frac{1}{p}\left(b - \frac{C}{1-C}\right) \in (0, 1)$.*

- *If $C \in [C_0, 1]$, D acquires information iff $k \leq k_0$; if $k = k_0$ this may be with probability $\gamma \in [0, 1]$. If uninformed, D chooses $a = 1$; if informed, D chooses $a = 1$ iff $\theta \geq \theta_0$ and $a = 0$ otherwise. M never investigates.*

Proof. See Appendix. ■

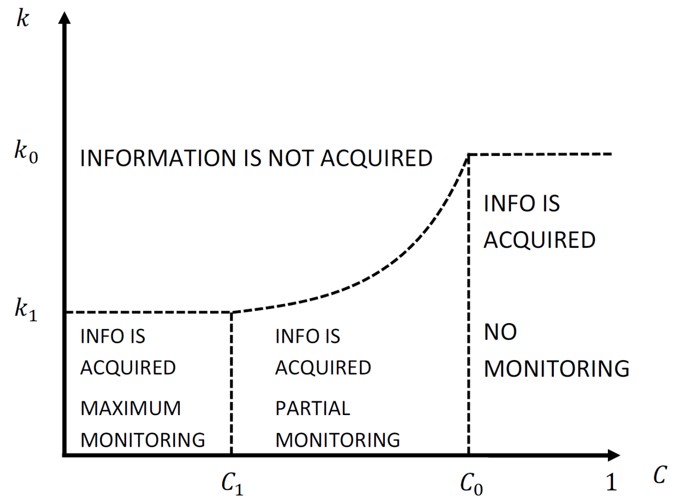

FIGURE 5: Equilibrium with observable information acquisition (Proposition 3)

Comparing Figure 5 with Figure 4, observable information acquisition affects only the region where, with unobservability, information is acquired with probability $\gamma \in (0,1)$. With observability no information is acquired in this region, heightening the chilling effect of freedom of information rules.

The intuition for this finding is as follows. When information acquisition occurs with some (intermediate) probability and is unobservable, when observing $a = 1$ the monitor cannot distinguish between an uninformed decision-maker, whom she would not want to investigate, and an informed one, whom she may wish to investigate. Uncertainly over whether information has been acquired reduces her incentive to investigate accordingly. With observable information acquisition the monitor can condition her investigation strategy on whether or not information has been acquired, increasing the probability that an informed decision-maker faces an investigation. Hence, unobservability partially protects the informed decision-maker from investigation, increasing his incentive to acquire information.

3.2 Voluntary disclosure

We now modify the monitoring game with unobservable information acquisition by adding the possibility that, assuming information has been acquired, D may voluntarily disclose θ at the same time as taking action a. We represent this as the decision-maker choosing $v \in \{0,1\}$ where $v = 1$ is voluntary disclosure and $v = 0$ is no voluntary disclosure.

We assume now that the decision-maker suffers a small cost ε of undergoing an investigation: even if investigation finds that he has taken the socially optimal action, the decision-maker nonetheless faces the inconvenience of responding to the monitor's request. Given this cost, an informed decision-maker may sometimes prefer to reveal his information rather than face an investigation. We assume ε to be negligible and generally omit it from calculations, but it appears in proofs and discussions whenever it breaks an indifference tie.

We assume that the monitor is sufficiently sophisticated that, if information is not disclosed, she can infer from the decision-maker's equilibrium strategy a posterior on θ, but that in order to generate reward r and punishment p she nonetheless needs to investigate to obtain evidence. Therefore, the monitor can use her inference only in deciding whether to investigate or not.

In this extension to the basic model, Lemma 6 still holds, i.e. $\underline{\pi} = 0$. However, we need additionally to determine the decision-maker's optimal disclosure strategy. The following lemma (which takes the place of Lemma 5 for the model without voluntary disclosure) describes the informed decision-maker's optimal choices regarding both action and disclosure.

Lemma 7 *For any $\{\underline{\pi}, \overline{\pi}\}$, an informed D chooses $a = 1$ iff his type is weakly larger than $\theta_{\overline{\pi}} \equiv \frac{1-b+\overline{\pi}p}{2}$ and $a = 0$ otherwise. Types $\theta \geq \frac{1}{2}$ strictly prefer to disclose their information iff $\overline{\pi} > 0$, otherwise they are indifferent towards disclosure. Types $\theta < \theta_{\overline{\pi}}$, who choose $a = 0$, are indifferent towards disclosure ($v \in \{0,1\}$).*

Proof. The optimal action is derived in a similar way to the proof of Lemma 5. For types $\theta \geq \frac{1}{2}$, the expected benefit from $a = 1$ is $b - 1 + 2\theta + \underline{\pi}p$ which is strictly positive. For types $\theta < \frac{1}{2}$, the expected benefit from $a = 1$ is $b - 1 + 2\theta - \overline{\pi}p$ which is weakly positive for $\theta \geq \theta_{\overline{\pi}} \equiv \frac{1-b+\overline{\pi}p}{2}$. The optimal disclosure strategy is derived by observing that types $\theta < \theta_{\overline{\pi}}$ and $\theta \geq \frac{1}{2}$ take the socially (as well as privately) optimal action, hence if the probability of undergoing an investigation is strictly positive then they strictly prefer to reveal their information, otherwise they are indifferent between disclosing and not disclosing. Since (from Lemma 6) $\underline{\pi} = 0$, types $\theta < \theta_{\overline{\pi}}$ face no threat of investigation and are indifferent towards disclosure. ■

The possibility of voluntary disclosure accompanying action $a = 1$ affects the equilibrium of the game as follows. Lemma 7 tells us that, when threatened with a strictly positive probability of investigation, informed types taking the socially optimal action choose to reveal their information. This restricts the ability of types that wish to act contrary to the social optimum to hide their misbehaviour from the monitor: if all others taking the relevant action reveal their information voluntarily, then the monitor will infer from non-disclosure that the decision-maker is taking the socially inferior action and will investigate accordingly.

The following proposition summarises the characterisation of equilibrium with voluntary disclosure. In Figure 6, which illustrates this proposition, disclosure strategies are noted for types choosing $a = 1$.

Proposition 4 (Equilibrium with voluntary disclosure)

- *If $k = k_0$ and $C \geq \gamma C_0$, there exists an equilibrium in which D acquires information with probability $\gamma \in (0,1)$. If uninformed, D chooses $a = 1$. If informed, D chooses $a = 1, v = 0$ iff $\theta \geq \theta_0$; otherwise D chooses $a = 0$ and $v \in \{0,1\}$. M never investigates.*

- *If $k < k_0$ and $C > C_0$, there exists an equilibrium in which D acquires information. D chooses $a = 1$ iff $\theta \geq \theta_0$ and $a = 0, v \in \{0,1\}$ otherwise. Types $\theta \in \left[\theta_0, \frac{1}{2}\right)$ choose $v = 0$. A proportion $\rho \in (0,1]$ of types $\theta \geq \frac{1}{2}$ choose $v = 0$. M never investigates.*

- *If $k \in (k_1, k_0)$ and $C \in (0,1)$, there exists an equilibrium in which D acquires information with probability $\gamma = C$. If uninformed, D chooses $a = 1$. If informed, D chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}_k}$; otherwise D chooses $a = 0$ and $v \in \{0,1\}$. Types $\theta \in \left[\frac{1}{2}, 1\right]$ choose $v = 1$, while types $\theta \in \left[\theta_{\overline{\pi}_k}, \frac{1}{2}\right)$ choose $v = 0$. After observing $a = 1$ and $v = 0$ M investigates with probability $\overline{\pi}_k = \frac{1}{p}\left(b - \sqrt{2b + 4k - 1}\right) \in (0,1)$.*

- *If $k \leq k_1$ and $C \in [0,1]$, there exists an equilibrium in which D acquires information. D chooses $a = 1, v = 1$ iff $\theta \geq \theta_1$ and $a = 0, v \in \{0,1\}$ otherwise. After observing $a = 1, v = 0$ the monitor's beliefs must be such that she investigates with probability one.*

- *If $k > k_0$, D never acquires information, and chooses $a = 1$. M never investigates.*

Proof. See Appendix. ■

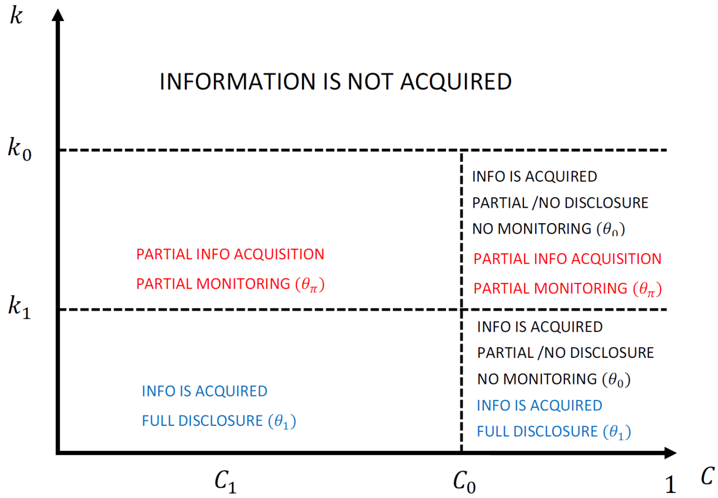

FIGURE 6: Equilibrium with voluntary disclosure (Proposition 4)

This model, unlike the previous cases, has multiple equilibria for $C > C_0$, $k < k_0$ (we adopt no refinements to select between these); below C_0 and above k_0 equilibrium is unique. As before, information is never acquired for $k > k_0$. The following discussion refers to disclosure *accompanying action* $a = 1$; as noted in Lemma 7, when taking $a = 0$ the decision-maker may either disclose or not, but this choice has no further implications.

For $C > C_0$ there is an equilibrium with full information acquisition and either partial or no disclosure which implements the private optimum (θ_0). Types $\theta \in \left[\theta_0, \frac{1}{2}\right)$, who act contrary to the social optimum, do not disclose; types $\theta \geq \frac{1}{2}$ (who choose $a = 1$ in line with the social optimum) may disclose or not but the proportion disclosing must be sufficiently low that the monitor has no incentive to investigate.

For $k < k_1$ and any C, there is an equilibrium in which information is acquired and fully disclosed; as a result the decision-maker's action choice is maximally disciplined (θ_1). In equilibrium no investigation occurs as information is always disclosed; however, full disclosure requires sufficiently sceptical off-equilibrium path beliefs that non-disclosure triggers an investigation.

For intermediate $k \in (k_1, k_0)$ and any C, there is a mixed strategy equilibrium in which information is acquired with probability $\gamma = C$; types taking $a = 1$ when this is the social optimum disclose while those taking this action contrary to the social optimum do not. If she observes $a = 1$ and no disclosure, the monitor investigates with probability $\overline{\pi}_k$, thus action is partially disciplined ($\theta_{\overline{\pi}_k}$). As before the probability that information is acquired increases with C, ensuring that the monitor remains indifferent between investigating and not investigating; similarly the probability of investigation is decreasing in k, ensuring the decision-maker's indifference towards information acquisition.

Comparing Proposition 4 with Proposition 2, voluntary disclosure extends the discipline effect of monitoring to higher values of C. However, the chilling effect on information acquisition is heightened too: a mixed equilibrium in which information is acquired with probability $\gamma \in (0, 1)$ now exists for parameter values where, in the absence of voluntary disclosure, information is acquired for sure. Moreover, comparing the mixed equilibria in Propositions 2 and 4, the probability of investigation after $a = 1$ (and no disclosure) is the same in the two cases – i.e. the discipline effect is unchanged – but with voluntary disclosure information is acquired with lower probability ($C < \gamma(k, C)$ for $k \in [k_1, k_0]$; the magnitude of γ is less than half as great) – in other words, the chilling effect is greater. Intuitively, the same $\overline{\pi}_k$ is required in the two cases in order that the decision-maker is indifferent towards information acquisition. However, the condition for the monitor to be indifferent towards investigation is now different. In Proposition 2 the monitor's expected reward from investigating depends on both the probability that the decision-maker is informed (i.e. γ) and the probability that an informed decision-maker takes the socially undesirable action. In the mixed equilibrium with voluntary disclosure, by contrast, all informed types taking the socially optimal action disclose their information, the second probability goes to unity and so, to generate the same expected reward, the probability of information acquisition must be lower.

The outcome of the model with voluntary disclosure may be understood by reference to the literature on persuasion: see Milgrom (1981), Milgrom (2008) and Milgrom and Roberts (1986). In common with papers on persuasion, and in contrast to those on 'cheap talk', in our analysis information is taken to be 'hard' and cannot be distorted. The persuasion literature typically finds that the ability to disclose information harms those that would wish to hide their type. However, our paper differs from the persuasion literature in two respects. First, the sender (here, the decision-maker) himself takes an action and so has a private value of information distinct from its value in communication with the receiver (here, the monitor). Secondly, the unraveling argument that typically guarantees full disclosure of information in persuasion games does not apply to our model. The unraveling argument requires that, in a candidate equilibrium with less-than-full disclosure, at least one type has an incentive to deviate and reveal his information in order to induce a more favourable response from the receiver. In our model, the receiver's response is the monitor's investigation decision , while the type(s) who might have an incentive to reveal their information are those who have taken the socially optimal action and wish to avoid the cost of undergoing an investigation. If the cost/reward ratio for investigation is sufficiently high, the threat of opening an investigation unless the decision maker discloses information is not credible, thus the unraveling argument does not apply.

4 Discussion and conclusion

Freedom of information legislation – which allows interested parties, journalists and individual citizens to access information held by public bodies – is often presented as an important measure to increase the transparency of decision-making by governments and other public bodies. The analysis in this paper demonstrates the role of freedom of information rules in disciplining the actions of decision-makers by providing access to such information. However, the analysis also highlights a potential drawback of such measures: when information is costly to acquire, freedom of information rules can also inhibit the acquisition of information, resulting in less informed decision-making.

In our framework, measures which strengthen freedom of information can be characterised as lowering the cost/reward ratio C for investigation of the decision-maker's action. The impact of reducing C depends critically on the cost of information acquisition, k. If k is sufficiently small, lowering C has no effect on information acquisition and raises the intensity of monitoring, increasing the discipline on the decision-maker's action. Discussions of the impact of freedom of information laws typically assume this is the situation: that increasing access to information improves transparency and discipline, without affecting the information acquired by public bodies.

However, our analysis shows that if information acquisition is more costly, freedom of information legislation can have unintended consequences. In our baseline model, reducing C below a critical level (so that the mixed equilibrium obtains) reduces information acquisition with no increase in monitoring intensity. Over this interval improving access to information worsens outcomes by reducing information-gathering by public bodies, resulting in less informed decisions. If the acquisition of information is observable to outsiders the situation is worse still: in this region information is not acquired at all and decisions are entirely uninformed.

Giving decision-makers the opportunity to disclose information voluntarily does not necessarily improve the situation. With voluntary disclosure the discipline effect of freedom of information is extended to higher cost-reward ratios: the tendency for information to be revealed makes it harder for a decision-maker that acts contrary to the public interest to hide this fact. But the chilling effect on information acquisition is heightened too: for intermediate levels of k there is only partial information acquisition, and the probability of information being acquired is less than half that found in the baseline model without voluntary disclosure.

These results have the following implications for freedom of information rules. If, generally, the kinds of information relevant to decision-making are freely available or cheap to acquire, then freedom of information rules are likely to have a beneficial effect in disciplining decision-making while having little impact on the informedness of public bodies. But when acquiring and processing information is costly, freedom of information rules may inhibit these investments. Moreover, improving outsiders' access to information, or making it easier for them to observe that information has been acquired, is likely to worsen this chilling effect.

In well-developed democratic systems, the chilling effect on information acquisition might be countered by rules and norms requiring public bodies to obtain and publish information when making their decisions. For example, in the UK investigations by competition authorities and decision-making by sectoral regulators (such as Ofcom, Ofgem, etc.) are now accompanied by several rounds of working papers and consultations, and information submitted by interested parties must also be taken into account. Such processes – underpinned by the threat of appeal if they are not followed – may mitigate the underinvestment in information highlighted in this paper.

This suggests the following lesson for institutional design. Freedom of information rules might be seen as a quick and easy means of improving oversight of public bodies – in effect, every citizen becomes a potential monitor. But in a less developed system, decision-makers may respond to the increased ease of access by reducing their efforts to acquire information, resulting in decisions being less informed. Unless requirements are also placed on public bodies to acquire and publish information to guide their decisions, the introduction of freedom of information legislation might worsen outcomes overall.

Appendix

Proof of Proposition 2

The probability with which the decision-maker acquires information is denoted γ. We consider in turn the three possibilities, characterising equilibrium in each case: $\gamma = 0$, $\gamma = 1$ and $\gamma \in (0, 1)$. Recall that, from Lemma 6, the monitor never investigates after observing $a = 0$ (i.e. it is always the case that $\underline{\pi} = 0$).

(i) Candidate equilibria with $\gamma = 0$

From Lemma 3, having acquired no information D optimally chooses $a = 1$, with expected payoff $b - \frac{1}{2}$. With equilibrium beliefs, Lemma 3 tells us that M optimally chooses not to investigate (i.e. $\overline{\pi} = 0$). Deviation by M to any $\overline{\pi} > 0$ is strictly unprofitable.

Suppose that D were to deviate by acquiring information. If he then chooses $a = 1$, the deviation is undetected by M, hence $\overline{\pi} = 0$. From Lemma 5, it follows that the informed D acts according to threshold θ_0. From (6), D's expected utility from acquiring information (also taking into account its cost) is $\frac{1}{4}\left(b^2 + 2b - 1\right) - k$. Hence, such a deviation is unprofitable iff $b - \frac{1}{2} \geq \frac{1}{4}\left(b^2 + 2b - 1\right) - k$, i.e. iff $k \geq k_0$.

(ii) Candidate equilibria with $\gamma = 1$

Suppose that $\overline{\pi} = 0$. From Lemma 5, D follows threshold θ_0 and receives expected payoff $\frac{1}{4}\left(b^2 + 2b - 1\right) - k$. First, we check that the monitor does not have a profitable deviation. M's expected payoff from investigation after observing $a = 1$ is

$$\int_{\frac{1-b}{2}}^{\frac{1}{2}} \frac{r}{1 - \frac{1-b}{2}} d\theta - c,$$

which is negative iff $\frac{c}{r} \geq C_0 \equiv \frac{b}{b+1}$. Secondly, we check that a deviation to $\gamma = 0$ is unprofitable for D. From Lemma 3, it follows that the expected payoff of such a deviation is $b - \frac{1}{2}$. Hence, it is unprofitable iff $k \leq k_0$.

Next, suppose that $\overline{\pi} = 1$. From Lemma 5, an informed decision maker chooses $a = 1$ iff $\theta \geq \theta_1$ and receives expected utility

$$\begin{aligned} EU &= \int_0^{\theta_1} -\theta d\theta + \int_{\theta_1}^{1} (b - (1-\theta))\, d\theta - \int_{\theta_1}^{\frac{1}{2}} p d\theta - k \\ &= \frac{1}{4}b^2 - \frac{1}{2}bp + \frac{1}{2}b + \frac{1}{4}p^2 - \frac{1}{4} - k. \end{aligned}$$

Given D's strategy, M's strategy is optimal iff the cost/reward ratio of investigating after observing $a = 1$ is sufficiently low. More precisely, this is the case iff

$$\int_{\theta_1}^{\frac{1}{2}} \frac{r}{1-\theta_1} d\theta - c \geq 0,$$

that is, if $C \leq C_1 \equiv \frac{b-p}{b-p+1}$. If D deviates to $\gamma = 0$, his expected payoff is $b - \frac{1}{2}$, hence the deviation is unprofitable iff

$$\frac{1}{4}b^2 - \frac{1}{2}bp + \frac{1}{2}b + \frac{1}{4}p^2 - \frac{1}{4} - k \geq b - \frac{1}{2},$$

i.e., iff $k \leq k_1 \equiv \frac{(1-b)^2 - p(2b-p)}{4} < k_0$. Hence we conclude that such an equilibrium exists for $C \leq C_1$ and $k \leq k_1$.

Finally, suppose that $\overline{\pi} \in (0, 1)$. From Lemma 5, D chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}}$. Hence, after observing $a = 1$ the monitor is willing to randomise iff

$$\int_{\theta_{\overline{\pi}}}^{\frac{1}{2}} \frac{r}{1-\theta_{\overline{\pi}}} d\theta = c,$$

i.e. iff $C = C_{\overline{\pi}} \equiv \frac{b-\overline{\pi}p}{b-\overline{\pi}p+1}$. Hence, for a given value of C, M's strategy must satisfy $\overline{\pi}_C = \frac{1}{p}\left(b - \frac{C}{1-C}\right)$, which in turn implies that the action threshold followed by an informed decision maker is $\theta_{\overline{\pi}(C)} = \frac{1-2C}{2(1-C)}$. The associated expected utility for D is given by

$$\int_0^{\theta_{\overline{\pi}_C}} (-\theta)\, d\theta + \int_{\theta_{\overline{\pi}_C}}^{1} (b-(1-\theta))\, d\theta - \int_{\theta_{\overline{\pi}_C}}^{\frac{1}{2}} \left(b - \frac{C}{(1-C)}\right) d\theta - k$$
$$= \frac{2C-1}{4(1-C)^2} + \frac{b}{2} - k.$$

Hence D has no incentive to deviate to $\gamma = 0$ iff $b - \frac{1}{2} \leq \frac{2C-1}{4(1-C)^2} + \frac{b}{2} - k$, i.e. for $k \leq k_C$ where

$$k_C \equiv \frac{1-b}{2} - \frac{1-2C}{4(1-C)^2}.$$

For this to be an equilibrium conditions are needed to ensure that $\overline{\pi}_C \in (0,1)$. Notice that $\overline{\pi}_C(C_1) = 1$ and $\overline{\pi}_C(C_0) = 0$. Hence, the above mixed strategy equilibrium exists for $C \in (C_1, C_0)$ and $k \leq k_C$. Also, notice that $k_C(C_1) = k_1$ and $k_C(C_0) = k_0$.

Candidate equilibria with $\gamma \in (0,1)$

Suppose that $\overline{\pi} = 0$. The above analysis for $\gamma = 1$ implies that, when informed, D acts according to threshold θ_0. For M optimally never to investigate requires the investigation cost weakly to exceed her expected reward; given the probability of information acquisition $\gamma \in (0,1)$ this is the case iff $C \geq \gamma C_0$. For D to be indifferent between acquiring and not acquiring information requires $k = k_0$. Hence we conclude that an equilibrium with $\gamma \in (0,1)$, $\overline{\pi} = 0$ exists iff $C \geq \gamma C_0$ and $k = k_0$.

Next, suppose that $\overline{\pi} = 1$. The above analysis for $\gamma = 1$ implies that, when informed, D acts according to threshold θ_1. For M optimally to investigate after $a = 1$ requires her expected reward weakly to exceed the investigation cost; given the probability of information acquisition $\gamma \in (0,1)$ this is the case iff $C \leq \gamma C_1$. For D to be indifferent between acquiring and not acquiring information requires $k = k_1$. Hence we conclude that an equilibrium with $\gamma \in (0,1)$, $\overline{\pi} = 1$ exists iff $C \leq \gamma C_1$ and $k = k_1$.

Finally, suppose that $\overline{\pi} \in (0,1)$. The above analysis for $\gamma = 1$ implies that, when informed, D acts according to threshold $\theta_{\overline{\pi}}$. For M to be indifferent between investigating and not investigating after observing $a = 1$ requires her expected reward from investigating to equal zero, i.e.

$$(1-\gamma)(-c)+\gamma\left(\int\limits_{\frac{1-b+\overline{\pi}p}{2}}^{\frac{1}{2}}(r-c)\,d\theta+\int\limits_{\frac{1}{2}}^{1}(-c)\,d\theta\right)=0,$$

which requires

$$C=\frac{\gamma(b-\overline{\pi}p)}{2+\gamma(b-\overline{\pi}p-1)}.$$

Solving for γ, we obtain

$$\gamma(C)=\frac{2C}{C+(1-C)(b-\overline{\pi}p)}.$$

For D to be indifferent at the information acquisition stage requires

$$b-\frac{1}{2}=\frac{(2b+b^2-1)-\overline{\pi}p(2b-p)}{4}-k.$$

Solving for $\overline{\pi}$ we obtain

$$\overline{\pi}_k=\frac{1}{p}\left(b-\sqrt{2b+4k-1}\right).$$

Observing that $\frac{\partial\overline{\pi}(k)}{\partial k}<0$ and imposing $\overline{\pi}(k)\in[0,1]$ we obtain the parameter condition $k\in[k_0,k_1]$. Substituting for $\overline{\pi}$ in the expression for γ, we obtain

$$\gamma(k,C)=\frac{2C}{C+(1-C)\left(\sqrt{2b+4k-1}\right)}.$$

Observing that $\frac{\partial\gamma(k,C)}{\partial k}<0$ and $\frac{\partial\gamma(k,C)}{\partial\frac{c}{r}}>0$ we obtain the condition $C\in\left(0,\frac{\sqrt{2b+4k-1}}{\sqrt{2b+4k-1}+1}\right)$. Notice that the locus $C=\frac{\sqrt{2b+4k-1}}{\sqrt{2b+4k-1}+1}$ corresponds to the locus $k=\frac{1-b}{2}-\frac{1-2C}{4(1-C)^2}$, which is $k(C)$. ■

Proof of Proposition 3

The proof follows very closely the proof of Proposition 2. Lemmas 3, 5 and 6 also hold in this special case; recall that Lemma 6 implies that when M observes that D is uninformed, she does not open an investigation. Again we consider in turn the three possible cases: $\gamma = 0$, $\gamma = 1$ and $\gamma \in (0, 1)$.

Candidate equilibria with $\gamma = 1$

The characterisation of equilibria with $\gamma = 1$ is identical to the one contained in the proof of Proposition 2. A deviation to $\gamma = 0$ is observable in this case, but this does not change the equilibrium in the ensuing subgame, in which D's expected utility is $b - \frac{1}{2}$.

Candidate equilibria with $\gamma = 0$

From Lemma 3, D chooses $a = 1$, with expected payoff $b - \frac{1}{2}$. Suppose that D deviates to acquire information. The deviation is observed by M. Three possible types of equilibria may occur in the ensuing subgame.

First, suppose that in this subgame $\overline{\pi} = 1$. Arguments analogous to those presented in the proof of Proposition 2 show that such an equilibrium exists iff $C \leq C_1$. The deviating D chooses $a = 1$ iff $\theta \geq \theta_1$ and achieves expected utility $\frac{1}{4}b^2 - \frac{1}{2}bp + \frac{1}{2}b + \frac{1}{4}p^2 - \frac{1}{4} - k$. Hence, an equilibrium with $\gamma = 0$ and such that $\overline{\pi} = 1$ (if D were to deviate to $\gamma = 1$) exists iff $C \leq C_1$ and $k \geq k_1$.

Next, suppose that in this subgame $\overline{\pi} = 0$. Arguments analogous to those presented in the proof of Proposition 2 show that such an equilibrium exists iff $C \geq C_0$. The deviating D chooses $a = 1$ iff $\theta \geq \theta_0$ and achieves expected utility $-\frac{1}{4}\left(1 - 2b - b^2\right) - k$. Hence, an equilibrium with $\gamma = 0$ and such that $\overline{\pi} = 0$ (if D were to deviate to $\gamma = 1$) exists iff $C \geq C_0$ and $k \geq k_0$.

Finally, suppose that in this subgame $\overline{\pi} \in (0, 1)$. Arguments analogous to those presented in the proof of Proposition 2 show that such an equilibrium exists iff $C \in (C_1, C_0)$. In this equilibrium, M investigates after $a = 1$ with probability $\overline{\pi}_C = \frac{1}{p}\left(b - \frac{C}{1-C}\right)$, D chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}(C)} = \frac{1-2C}{2(1-C)}$ and achieves expected utility $\frac{2C-1}{4(1-C)^2} + \frac{b}{2} - k$. Hence, an equilibrium with $\gamma = 0$ and such that $\overline{\pi} = 0$ (if D were to deviate to $\gamma = 1$) exists iff $k \geq k_C = \frac{1-b}{2} - \frac{1-2C}{4(1-C)^2}$ and $C \in (C_1, C_0)$.

Candidate equilibria with $\gamma \in (0,1)$

It follows immediately from the above analysis that these equilibria may occur only as the limiting case of the equilibria with $\gamma \in \{0,1\}$. More precisely, for $C \leq C_1$ and $k = k_1$ there exists an equilibrium in which D randomises at the information acquisition stage; if information is acquired, D chooses $a = 1$ iff $\theta \geq \theta_1$ and M investigates after observing $a = 1$.

For $C \in (C_1, C_0)$ and $k = k_C$ there exists an equilibrium in which D randomises at the information acquisition stage; if information is acquired, D chooses $a = 1$ iff $\theta \geq \frac{1-2C}{2(1-C)}$ and after observing $a = 1$ M investigates with probability $\overline{\pi}_C = \frac{1}{p}\left(b - \frac{C}{1-C}\right)$. ■

Proof of Proposition 4

We start by observing that Lemma 3 also holds in this case. The optimal choice of an informed decision maker is now described by Lemma 7 (rather than Lemma 5).

Recall that, from Lemma 6, the monitor never investigates after observing $a = 0$ (i.e. $\underline{\pi} = 0$). Accordingly, Lemma 7 tells us that when D chooses $a = 0$, he is indifferent towards disclosure (i.e. $v \in \{0,1\}$). The following establishes D's equilibrium disclosure choices when choosing $a = 1$ and M's equilibrium investigation strategies after observing $a = 1$. We consider in turn the three possible cases for information acquisition: $\gamma = 0$, $\gamma = 1$ and $\gamma \in (0,1)$.

Candidate equilibria with $\gamma = 0$

From Lemma 3, it follows that in equilibrium M does not investigate (i.e. $\overline{\pi} = 0$) and D chooses $a = 1$, gaining expected utility $b - \frac{1}{2}$. Suppose D deviates to $\gamma = 1$. It follows from Lemma 7 that he chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}}$. Since a deviation to $\gamma = 1, a = 1$ is undetectable by M (unless D reveals his information, in which case M does not need to investigate anyway), $\overline{\pi} = 0$. Hence D's action threshold is θ_0 and all types $\theta \in \left[\theta_0, \frac{1}{2}\right)$ strictly prefer not to disclose their information in order to avoid punishment. This implies that for any off-equilibrium path beliefs following $a = 0$, hence for any $\underline{\pi}$, the highest payoff achievable by deviating and acquiring information is $\frac{1}{4}(b^2 + 2b - 1) - k$. Hence, we can conclude that with any off-equilibrium path beliefs held by M, an equilibrium with $\gamma = 0$ exists iff $k \geq k_0$.

Candidate equilibria with $\gamma = 1$

For simplicity, we describe candidate equilibria in three sets depending on the equilibrium disclosure choice of the types choosing $a = 1$.

Equilibria with $\gamma = 1$ *and full disclosure* ($v = 1$) *accompanying* $a = 1$.

Consider the informed D's action choice. For types $\theta \geq \frac{1}{2}$, the relative benefit from choosing $a = 1$ (as compared with $a = 0$) is $b - 1 + 2\theta$, which is strictly positive. For types $\theta < \frac{1}{2}$, the relative benefit from $a = 1, v = 1$ (as compared with $a = 0$) is $b - 1 + 2\theta - p$, which is weakly positive for $\theta \geq \theta_1$. Accordingly, in an equilibrium where $a = 1$ is accompanied by full disclosure, D chooses $a = 1$ iff $\theta \geq \theta_1$. D's equilibrium expected payoff is $\frac{1}{4}b^2 - \frac{1}{2}bp + \frac{1}{2}b + \frac{1}{4}p^2 - \frac{1}{4} - k$.

Now consider the possible deviations. Types $\theta \geq \frac{1}{2}$ achieve their maximum payoff conditional on having acquired information, so have no incentive to deviate to no disclosure and either action. Types in $\left[\theta_0, \frac{1}{2}\right)$ may have an incentive to deviate to $a = 1, v = 0$ to reduce the probability of punishment: this deviation is unprofitable for all of these types iff $\overline{\pi} = 1$, which in turn requires that M's off-equilibrium path beliefs after observing $a = 1, v = 0$ assign a sufficiently high probability to the case $\gamma = 1, \theta < \frac{1}{2}$, compared to the value of C, so as to guarantee that opening an investigation is optimal. (If, after observing a deviation to $a = 1, v = 0$, M believes that $\gamma = 1, \theta < \frac{1}{2}$ for sure then $\overline{\pi} = 1$ is optimal for any C.)

Finally, consider a deviation to $\gamma = 0$. The deviation payoff is $b - \frac{1}{2}$, compared with equilibrium payoff $\frac{1}{4}b^2 - \frac{1}{2}bp + \frac{1}{2}b + \frac{1}{4}p^2 - \frac{1}{4} - k$, thus such an equilibrium exists iff $k \leq k_1$.

Equilibria with $\gamma = 1$ *and no disclosure* ($v = 0$) *accompanying* $a = 1$.

Consider the informed D's action choice. From Lemma 7, D chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}}$. For $a = 1, v = 0$ to be an equilibrium it must be the case that $\overline{\pi} = 0$: for any $\overline{\pi} > 0$ types $\theta \geq \frac{1}{2}$ would profitably deviate to $v = 1$. This implies that

$\theta_{\overline{\pi}} = \theta_0$. Types $\theta \in \left[\theta_0, \frac{1}{2}\right)$ have no incentive to deviate to $v = 1$, as this would imply a punishment. For M, $\overline{\pi} = 0$ is compatible with equilibrium only if $C \geq C_0$.

Finally, consider a deviation to $\gamma = 0$. Such a deviation is unprofitable iff $k \leq k_0$, as in equilibrium D achieves his privately optimal payoff.

Equilibria with $\gamma = 1$ *and partial disclosure* ($v \in \{0, 1\}$) *accompanying* $a = 1$.

Consider the informed D's action choice. From Lemma 7, D chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}}$. First, we show that $\overline{\pi} = 0$. Suppose instead that $\overline{\pi} > 0$: in equilibrium all types $\theta \geq \frac{1}{2}$ would disclose their type to avoid the threat of investigation, hence all non-disclosing types would have $\theta < \frac{1}{2}$. M would then optimally choose $\overline{\pi} = 1$ (for any C). But then all types choosing $a = 1$ would have a strict incentive to disclose: even those taking the socially inferior action would disclose in order to avoid the inconvenience of being investigated, ε. Thus partial disclosure is possible only if $\overline{\pi} = 0$. Having established this, it follows that D chooses $a = 1$ iff $\theta \geq \theta_0$.

Denote by s_1 the measure of the set of types who choose $a = 1$ in equilibrium, and by $\underline{s_1}$ the measure of the subset of s_1 composed of types $\theta < \frac{1}{2}$. $\overline{\pi} = 0$ is compatible with equilibrium iff $C \geq \frac{\underline{s_1}}{s_1}$, the probability that investigation after observing $a = 1$ results in a reward. Notice that as the proportion of types $\theta \geq \frac{1}{2}$ that do not disclose converges to 1 this cost threshold converges to C_0; as the proportion of types $\theta \geq \frac{1}{2}$ that do not disclose converges to 0 this cost threshold converges to 1.

Finally, consider a deviation to $\gamma = 0$. Such a deviation is unprofitable iff $k \leq k_0$, as in equilibrium D achieves his privately optimal payoff.

Candidate equilibria with $\gamma \in (0, 1)$

As above, we describe candidate equilibria in three sets depending on the equilibrium disclosure choice of the types choosing $a = 1$.

Equilibria with $\gamma \in (0, 1)$ *and full disclosure* ($v = 1$) *accompanying* $a = 1$.

Consider the informed D's action choice. From Lemma 7, D chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}}$. It cannot be the case that $\overline{\pi} = 0$, because types $\theta \in \left(\theta_0, \frac{1}{2}\right)$ could then profitably deviate to $v = 0$. Also, from Lemma 3, it cannot be the case that $\overline{\pi} > 0$

because after observing $a = 1, v = 0$ the monitor, who holds equilibrium beliefs, expects D to be uninformed and would optimally choose not to investigate. Hence there exists no equilibrium in this class.

Equilibria with $\gamma \in (0,1)$ and no disclosure ($v = 0$) accompanying $a = 1$.

From Lemma 7, if D acquires information, he chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}}$. In equilibrium, types $\theta \geq \frac{1}{2}$ choose not to disclose only if $\overline{\pi} = 0$, otherwise they could profitably deviate to $v = 1$. Hence, it must be the case that $\overline{\pi} = 0$; accordingly an informed D chooses $a = 1$ iff $\theta \geq \theta_0$ and achieves expected payoff $\frac{1}{4}(b^2 + 2b - 1) - k$. After observing $a = 1, v = 0$, M, who holds equilibrium beliefs, expects investigation to yield a reward with probability $\gamma \left(\frac{\frac{1}{2} - \theta_0}{1 - \theta_0} \right) = \gamma C_0$. Hence, $\overline{\pi} = 0$ is her optimal choice iff $C \geq \gamma C_0$. From Lemma 3, D's payoff if he does not acquire information is $b - \frac{1}{2}$. Hence, $\gamma \in (0,1)$ is compatible with equilibrium iff $k = k_0$.

Equilibria with $\gamma \in (0,1)$ and partial disclosure ($v \in \{0,1\}$) accompanying $a = 1$.

From Lemma 7, if D acquires information, he chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}}$. There is no equilibrium in this class with $\overline{\pi} = 1$, because all types who choose $a = 1$ would then have a strict incentive to disclose their type and avoid the inconvenience of being investigated, ε.

Suppose that $\overline{\pi} = 0$. If D acquires information, he acts iff $\theta \geq \theta_0$. No type $\theta \in \left[\theta_0, \frac{1}{2}\right)$ discloses his information. Some types $\theta \in \left[\frac{1}{2}, 1\right]$ disclose. Given M's equilibrium beliefs, the condition for $\overline{\pi} = 0$ to be optimal is that $C \geq \gamma \frac{\underline{s}_1}{s_1} \in [\gamma C_0, \gamma)$ where the exact value of $\frac{\underline{s}_1}{s_1}$ depends on the disclosure choice of the types $\theta \in \left[\frac{1}{2}, 1\right]$. The condition for information acquisition to occur with probability $\gamma \in (0,1)$ is $k = k_0$.

Finally, suppose that $\overline{\pi} \in (0,1)$. From Lemma 7, if D acquires information, he chooses $a = 1$ iff $\theta \geq \theta_{\overline{\pi}}$. All types $\theta \in \left[\frac{1}{2}, 1\right]$ choose $v = 1$ to avoid an investigation, while all types $\theta \in \left[\theta_0, \frac{1}{2}\right)$ choose $v = 0$ in the hope of avoiding punishment (given there is a positive probability of no investigation). After observing $a = 1, v = 0$, M, who holds equilibrium beliefs, infers that with probability $(1 - \gamma)$ D is uninformed, and with probability γ D is informed and has a type smaller than one half. Therefore, M is indifferent between investigating and not investigation iff $c = \gamma r$, i.e. iff $C = \gamma$.

$\overline{\pi}$ is determined by the indifference condition for D at the information acquisition stage:

$$b - \frac{1}{2} = \int_0^{\frac{1-b+\overline{\pi}p}{2}} -\theta d\theta + \int_{\frac{1-b+\overline{\pi}p}{2}}^{\frac{1}{2}} (b - 1 + \theta - \overline{\pi}p)\, d\theta + \int_{\frac{1}{2}}^{1} (b - 1 + \theta)\, d\theta - k,$$

which gives us

$$k = \frac{(1-b)^2}{4} - \frac{\overline{\pi}p\,(2b-p)}{4}.$$

Solving for $\overline{\pi}$, we obtain

$$\overline{\pi}_k = \frac{1}{p}\left(b - \sqrt{2b + 4k - 1}\right).$$

We need to check that this expression is indeed a probability. First, notice that $\overline{\pi}_k$ is decreasing in k: for D to be indifferent between acquiring information or not, the higher information cost must be balanced by a lower probability of investigation. Evaluating the expression at $k = k_0$ and $k = k_1$ we obtain $\overline{\pi}_k\,(k_0) = 0$ and $\overline{\pi}_k\,(k_1) = 1$. Hence such an equilibrium exists for $k \in (k_1, k_0)$.

References

[1] Che, Yeon-Koo, and Navin Kartik (2009), "Opinions as incentives," *Journal of Political Economy*, 117(5), 815-860.

[2] Felgenhauer, Mike, and Petra Loerke (2013), "Public versus private experimentation," working paper.

[3] Felgenhauer, Mike, and Petra Loerke (2014), "Bayesian persuasion with private experimentation," working paper.

[4] Felgenhauer, Mike, and Elisabeth Schulte (2014), "Strategic private experimentation," *American Economic Journal: Microeconomics*, 6(4), 74-105.

[5] Frankel, Maurice (2001), "Freedom of information: Some international characteristics," Campaign for Freedom of Information.

[6] Hagenbach, Jeanne, Frederic Koessler and Eduardo Perez-Richet (2014), "Certifiable pre-play communication: Full disclosure," *Econometrica*, 82(3), 1093-1131.

[7] Hansen, Stephen, Michael McMahon and Andrea Prat (2014), "Transparency and deliberation within the FOMC: A computational lingustics approach," working paper.

[8] Hazell, Robert, Gabrielle Bourke and Benjamin Worthy (2012), "Open house? Freedom of information and its impact on the UK parliament," *Public Administration*, 90(4), 901-921.

[9] Hazell, Robert, and Mark Glover (2011), "The impact of freedom of information on Whitehall," *Public Administration*, 89(4), 1664-1681.

[10] Holmström, Bengt (1979), "Managerial incentive problems: A dynamic perspective," *Review of Economic Studies*, 6(1), 169-182.

[11] Hoppe, Eve (2013), "Observability of information acquisition in agency models," *Economic Letters*, 19, 104-107.

[12] Kamenica, Emir, and Matthew Gentzkow (2011), "Bayesian persuasion", *American Economic Review*, 101(6), 2590-2615.

[13] Kamenica, Emir, and Matthew Gentzkow (2014), "Disclosure of endogenous information," working paper.

[14] Levy, Gilat (2007), "Decision making in committees: Transparency, reputation, and voting rules," *American Economic Review*, 97(1), 150-168.

[15] Maskin, Eric, and Jean Tirole (2004), "The politician and the judge: Accountability in government," *American Economic Review*, 94(4), 1034-1054.

[16] Mathis, Jerome (2008), "Full revelation of information in Sender-Receiver games of persuasion," *Journal of Economic Theory*, 143, 571-84.

[17] Meade, Ellen E., and David Stasavage (2008), "Publicity of debate and the incentive to dissent: Evidence from the US Federal Reserve," *The Economic Journal*, 118 (April), 695–717.

[18] Milgrom, Paul (1981), "Good news and bad news: Representation theorem and applications," *Bell Journal of Economics*, 12, 380-91.

[19] Milgrom, Paul (2008), "What the seller won't tell you: Persuasion and disclosure in markets", *Journal of Economic Perspectives*, 22(2), 115-131.

[20] Milgrom, Paul, and John Roberts (1986), "Relying on the information of interested parties", *Rand Journal of Economics*, 17, 18-32.

[21] Prat, Andrea (2005), "The wrong kind of transparency," *American Economic Review*, 95(3), 862-877.

[22] Rayo, Luis, and Ilya Segal (2010), "Optimal information disclosure," *Journal of Political Economy*, 118(5), 949-987.

[23] Worthy, Benjamin, and Robert Hazell (2013), "The impact of the Freedom of Information Act in the UK," in Nigel Bowles, James T Hamilton and David

Levy (eds.), *Transparency in Politics and the Media: Accountability and Open Government*, London: L.B. Tauris, 31-45.