



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

On the (Mis-)Use of Information for Public Debate

Citation for published version:

Vikander, N & Patacconi, A 2013 'On the (Mis-)Use of Information for Public Debate'.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



On the (Mis-)Use of Information for Public Debate

Andrea Pataconi* Nick Vikander†

February 2013

Abstract

Policymakers often motivate their decisions by disclosing information. While this can help hold the government to account, it may also give policymakers an incentive to "fix the evidence" around their preferred policy. This paper considers a model of biased information gathering where the government can influence the workings of an agency in charge of collecting information. We examine how different disclosure rules and the degree of independence of the government agency affect citizen welfare. Our main result is that insulating the agency from political pressure, so that its information is always unbiased, may not be socially optimal. A biased information gathering process can curb the government's tendency to implement its ex ante favored policy, thus mitigating the agency conflict between policymakers and the public.

Keywords: Transparency, Accountability, Independence, Manipulation of Information.

JEL Classification: D73, H11, H56.

*University of Aberdeen Business School, Edward Wright Building, Dunbar Street, AB24 3QY, Aberdeen (UK). Email: a.pataconi@abdn.ac.uk

†University of Edinburgh, 30 Buccleuch Place, EH8 9JT, Edinburgh (UK). Email: nick.vikander@ed.ac.uk

Human experience teaches us that those who expect public dissemination of their remarks may well temper candor with a concern for appearances and for their own interest to the detriment of the decision-making process. (U.S. Supreme Court, *United States v. Nixon*)

We also recognize that there is a real dilemma between giving the public an authoritative account of the intelligence picture and protecting the objectivity of the JIC [Joint Intelligence Committee] from the pressures imposed by providing information for public debate. (Butler Report, p. 114)

1. Introduction

Transparency is an essential feature of a democratic and accountable state and yet, despite substantial progress in recent years, exceptions to the principle of open government remain commonplace (Prat, 2006). In the United States, for instance, the President has the right to withhold information from Congress and the courts, typically on the grounds that he needs candid and confidential advice from his staff. Freedom of Information laws also frequently allow policymakers to withhold information, most notably to protect internal decision making, personal privacy and national security (Banisar, 2004; Roberts, 2006).

This paper examines one important rationale for lack of transparency in government: the concern that public dissemination of information might compromise the quality of government decision making. We develop a model where the government receives information from an agency about a particular policy, and then decides whether or not the policy should be implemented. For instance, the government might receive an intelligence report about the opportunity to go to war, or an environmental impact assessment about the opportunity to build a new nuclear power plant. As is standard in political agency models, the preferences of the government and the public are not perfectly aligned. The government is more favorable than the public towards implementation but also wants public support for its decision. Thus, while policymakers may be more willing to wage war than voters, they are nevertheless responsive to public opinion.

Our key assumption is that the agency may be politicized and hence its report to the government may be biased. If the agency is independent, then it provides an unbiased report about the consequences of implementing the policy, and hence about the appropriate course of action. However, if the agency is *not* independent, then this report may be biased in favor of the government's ex ante preferred decision; that is, the report may be biased in favor of implementation. With a nonindependent agency, we assume that the government can choose

the optimal degree of bias so as to maximize its own welfare. For instance, the government may staff the agency with individuals who are prone to stating a case for war, seek the advice of biased experts, or encourage biased information gathering and evaluation. The drawback is that all parties with access to the report (including the government) then receive lower quality information which can result in poor decision making.¹

We use this framework to address two questions, both from the perspective of the public. First, should the contents of the report be publicly disclosed? And second, should the agency be made independent of the government? Both issues are of great practical importance. It is often claimed that secrecy is instrumental in protecting the integrity of the decision-making process and indeed one of the most common exemptions to the principle of open government concerns pre-decision information (Banisar, 2004). Granting independence to government agencies is also becoming increasingly common. The Federal Trade Commission in the U.S. and the Bank of England, for instance, have a status that ensures their independence from political pressure by limiting the removal of their heads to certain specific causes. The British commission in charge of investigating recent episodes of intelligence failure also recommended to strengthen the independence of the Joint Intelligence Committee, although it fell well short of recommending full independence from the executive (Butler Report, 2004, pp. 143-144).

In line with conventional wisdom, we find that disclosure ('transparency') makes the government more accountable and hence more responsive to public desires, relative to nondisclosure ('secrecy'). However, disclosure also induces policymakers to distort the process of information gathering and evaluation. In contrast, when no information can be disclosed, the government has no incentive to manipulate information. Secrecy is therefore effective at protecting the integrity of the decision-making process.

¹The recent debate on intelligence failures provides several examples of potentially biased information gathering. Consider the case of Curve Ball. Curve Ball was the codename of an Iraqi informant whose "revelations" constituted the backbone of the intelligence on Iraq's mobile biological weapons program during the run-up to the second Iraqi war. These "revelations", however, were later revealed to be complete fabrications. The most disturbing aspect of the Curve Ball fiasco is that concerns about the reliability of this informant appear to have been systematically suppressed. The Select Committee on Intelligence (2004), for instance, reports that when the CIA agent who had interviewed Curve Ball raised concerns about his reliability, he was told by the Deputy Chief of the CIA's Iraqi Task Force: "As I said last night, let's keep in mind the fact that this war's going to happen regardless of what Curve Ball said or didn't say, and that the Powers That Be probably aren't terribly interested in whether Curve Ball knows what he's talking about" (extract from an e-mail provided to the Committee, p. 249). The WMD Commission (2005) also notes that "the analysts who raised concerns about the need for reassessments were not rewarded for having done so but were instead forced to leave WINPAC" (p. 193). It seems that, as a result of this biased vetting process, relevant information was not transmitted to policymakers, thus potentially contributing to poor decision making.

We also consider a constitutional stage in which both the disclosure rule and the agency's degree of independence can be specified. The most surprising results emerge regarding what rule and degree of independence maximize the public's welfare. We show that from the public's perspective, secrecy is *never* optimal, but it can be optimal for the government agency *not* to be independent. Secrecy is always dominated by transparency because its chief advantage – unbiased information – can be more efficiently obtained by insulating the agency from political pressure. And yet the public may sometimes prefer that the agency be politicized so that its report is potentially biased. For any given decision rule, biased information increases the probability that the government will make the wrong decision, which hurts the public. However, the government wants to avoid making the wrong decision, so it tailors its optimal decision rule to the agency's level of bias. We show that a pro-implementation bias in information has a *moderating* effect; for given evidence, it makes the government more reluctant to implement the policy. This moderating effect benefits the public, which views implementation less favorably than the government. Thus, manipulation of information can help mitigate the agency conflict between the government and the public.

From a theoretical perspective, this result can be seen as an application of the theory of the second-best. According to this theory, introducing a new inefficiency – manipulation of information – in an environment where another inefficiency is already present – the agency conflict between the government and the public – can sometimes increase social welfare.

Previous work has examined how politicians can be held accountable when voters are not perfectly informed. Canes-Wrone et al. (2001) and Maskin and Tirole (2004) study models where policymakers have private information and reelection concerns create incentives for pandering. However, because these models do not allow policymakers to credibly communicate their private information to voters, they cannot distinguish between transparency and secrecy. Ashworth and Shotts (2010), building on Canes-Wrone et al. (2001), examine the effect of media bias. They find that pandering incentives can be lower when the media sometimes refrain from criticizing the government because negative media reports then become strongly indicative of an incorrect policy choice. Besley and Prat (2006) develop a model where incumbents (good and bad) can manipulate media reports by offering some form of compensation to the media owners. Their analysis focuses on how features of the media industry affect the quality of the media reports and political turnover. In contrast, we focus on how the integrity of the decision-making process can be protected when manipulation

incentives are present.²

This paper is also related to the literature on transparency in principal-agent relationships (e.g., Holmstrom, 1979; Cremer, 1995; Stasavage 2004). Prat (2005), in particular, develops a model of career concerns for experts where the principal can observe the agent's action and/or its consequences. He shows that transparency on action can induce the agent to disregard useful private information and act in a conformist manner. As a consequence, the principal can be better off by committing not to observe the action. Transparency on consequences, by contrast, always benefits the principal. Our focus is neither on transparency on action nor on consequences. We measure transparency by the extent to which *pre-decision information* is shared between the agent and the principal. Our focus is not whether transparency induces conformism on the part of the agent, but whether an agent will distort his own information (and possibly the principal's) to influence how the principal perceives his action.³

The remainder of the paper is organized as follows. The next section introduces the model. Sections 3 and 4 study different disclosure rules (transparency and secrecy), under the assumption that the government agency is nonindependent. Section 5 considers the case of an independent agency and compares different institutional arrangements from the public's point of view. Extensions are discussed in Section 6, while Section 7 concludes. Proofs are gathered in two technical appendices.

2. Model

We consider a model of government decision making where (i) the government is responsive to public opinion and (ii) the agency that provides the government with information is potentially biased. The model has four stages. At stage 1, if the agency is nonindependent, then the government chooses the agency's level of bias, $q \in [0, 1]$. One can interpret q as the type of bureaucrats who work at the agency. In contrast, if the agency is independent, its bias is equal to zero ($q \equiv 0$). At stage 2, the agency produces a report for the government. This report may or may not be publicly revealed, depending on the disclosure rule, as discussed

²See also Biglaiser and Mezzetti (1997), Milbourn et al. (2001) and Suurmond et al. (2004). These papers examine how reputational concerns affect the incentives to gather information and implement policies.

³Levy (2007) and Swank et al. (2008) develop models closely related to Prat's to study the effect of transparency on committee decision making. They show that secrecy can be conducive to better decision making because, if individual votes cannot be observed, then voters have less of an incentive to distort their actions in order to signal their types.

below. At stage 3, the government makes a policy decision $p \in \{a, n\}$, where a stands for implementation and n for the status quo. The public then decides whether to support the government's decision, $v \in \{a, n\}$. For example, if $p = a$, we will say that the government implements the policy, or that it selects implementation. If $v = a$, we will say that the public supports implementation. At stage 4, payoffs are realized.

Preferences. The payoffs of the government and the public depend on the state of the world, $S \in \{A, N\}$. For simplicity, we assume that A and N are a priori equally likely. The public would like the policy decision to match the true state, $a = A$ or $n = N$, in which case its payoff is zero. The public incurs a loss of C_a if the policy is implemented and the true state is N , and a loss of C_n if the policy is not implemented and the true state is A . Without loss of generality, $C_a + C_n = 1$.

The public supports the policy decision that maximizes its expected payoff. Let σ_P denote the public's posterior belief that the true state is A . The public supports implementation if it offers a higher expected payoff than the status quo,

$$-C_a(1 - \sigma_P) \geq -C_n\sigma_P \quad \Leftrightarrow \quad \sigma_P \geq \frac{C_a}{C_a + C_n} = C_a. \quad (2.1)$$

Thus, the public tends to support implementation if it believes state A is likely (σ_P high) and if the cost of mistaken implementation, C_a , is relatively small. For now, we arbitrarily break ties in favor of implementation.

The preferences of the government are more complex. First, the government cares about public welfare (a 'legacy' concern). It incurs a loss of C_a from implementing the policy when the state is N , and a loss of C_n from not implementing the policy when the state is A . In addition to this legacy concern, the government enjoys a private benefit $B \geq 0$ from implementation. Finally, the government incurs a loss $E \geq 0$ whenever its decision is not supported by the public. This cost captures in a stylized fashion the disciplining effect of public opinion. One interpretation of E is in terms of electoral concerns: if the government adopts an unpopular policy, citizens may vote for the opposition in the next election. However, E could also measure other costs associated with a loss of popularity, such as vilification by the press or reduced job opportunities in the private sector. The notion of "public" should also be interpreted broadly. For instance, the model could apply to settings where one country wants to convince others of a particular course of action to receive logistic or military support. In that case E would measure the loss to that country

should support be denied.

Let σ_{Gov} denote the government's belief that $S = A$. Let $1_{\{p,v\}}$ be an indicator function that takes a value of 1 if and only if a particular policy decision p is not supported by the public ($p \neq v$). Then the government chooses implementation over the status quo if

$$\sigma_{Gov} \geq C_a - B + (1_{\{a,v\}} - 1_{\{n,v\}})E. \quad (2.2)$$

Comparing (2.2) with (2.1), the government views implementation more favorably than the public due to the private benefit B , which creates a potential conflict of interest. However, the government can also be swayed by public opinion, as captured by the term $(1_{\{a,v\}} - 1_{\{n,v\}})E$. For instance, the government is more likely to implement the policy if the public supports implementation, in which case $v = a$ and $(1_{\{a,v\}} - 1_{\{n,v\}})E = -E$.

Information Structure. Before making a policy decision, the government receives a report from the agency. This report is composed of two signals, $s_i \in \{\alpha, \emptyset\}$, $i = 1, 2$. A α signal provides evidence in support of implementation, while a \emptyset signal provides evidence in support of the status quo. If the agency is independent, then these signals are genuine, s_i^G . Genuine signals are informative, conditionally independent and satisfy $\Pr(s_i^G = \alpha|A) = \Pr(s_i^G = \emptyset|N) = \theta$, where $\theta \in (\frac{1}{2}, 1)$ measures the signal precision.⁴

If the agency is nonindependent, then the signal-generating process may be distorted. Let $\mathbf{s}^q = \{s_1^q, s_2^q\}$ be the report produced by a nonindependent agency with bias $q \in [0, 1]$. We capture the idea of *asymmetric vetting* by assuming that with probability q , a genuine \emptyset signal is transformed into a fake α signal. That is, the nonindependent agency garbles the signal-generating process so that $\Pr(s_i^q = \alpha \mid s_i^G = \emptyset) = q$, which is independent across signals. The probability that a genuine α signal is transformed into a fake \emptyset signal is zero, $\Pr(s_i^q = \emptyset \mid s_i^G = \alpha) = 0$. Thus q measures the agency's bias in favor of the government's ex ante preference for implementation. A non-independent agency with zero bias will behave just like an independent agency, and produce a report consisting of genuine signals, $\mathbf{s}^0 = \mathbf{s}^G = \{s_1^G, s_2^G\}$.

If the agency is nonindependent, we allow the government to choose q to maximize its own payoff. This assumption is plausible if the government can appoint key agency person-

⁴We use two signals to allow for situations where the evidence is mixed. We use binary signals (instead of a single signal with multiple signal realizations) because this allows for a simple parametrization of the process of information manipulation (see below).

nel or can punish or reward them.⁵ We also posit that the government only observes the biased signals \mathbf{s}^q , rather than the genuine ones. This captures the fundamental drawback of manipulations: information is lost which may have been useful for decision making.⁶

Before proceeding, we introduce some additional notation. Let $(\cdot, \cdot)^q$ be a shorthand for $\mathbf{s}^q = (\cdot, \cdot)$. Any party that observes \mathbf{s}^q will update its beliefs about the true state, where we define

$$\sigma_+^q \equiv \Pr(A | (\alpha, \alpha)^q),$$

$$\sigma^q \equiv \Pr(A | (\alpha, \emptyset)^q) = \Pr(A | (\emptyset, \alpha)^q),$$

$$\sigma_- \equiv \sigma_-^q \equiv \Pr(A | (\emptyset, \emptyset)^q) = \sigma_-^G.$$

These beliefs correspond to the three possible cases that can arise: (i) the report supports implementation ($\mathbf{s}^q = (\alpha, \alpha)$), (ii) the report is mixed ($\mathbf{s}^q = (\alpha, \emptyset)$ or (\emptyset, α)) or (iii) the report supports the status quo ($\mathbf{s}^q = (\emptyset, \emptyset)$). We sometimes refer to α signals as positive signals, and to \emptyset signals as negative signals. Note that σ_- does not depend on q , because negative signals must be genuine.⁷

We also make the following assumptions:

Assumption 1. $C_a \in (\frac{1}{2}, \sigma_+^G]$,

Assumption 2. $B < C_a - \sigma_-$.

Assumption 1 states that the public always supports the status quo when the evidence is mixed, but would support implementation if it observed two positive, genuine signals. Assumption 2 can be interpreted as a weak form of congruency between the government and the public. It states that, even absent electoral concern ($E = 0$), the government's optimal

⁵Our analysis leaves the motivations of agency bureaucrats in the background, allowing us to focus on how disclosure rules and public opinion can shape policy. For theoretical analyzes of bureaucratic behavior, see for instance Prendergast (1993) or Alesina and Tabellini (2007, 2008).

⁶The assumption that the government only observes the biased signals \mathbf{s}^q implies that a government which discloses forged information is not telling a lie. In this sense, disclosure is truthful. Also note that the public does not explicitly penalize the government for manipulating information. This assumption could easily be relaxed by assuming the government incurs an additional cost $C(q)$ that is increasing in q .

⁷Simple computations yield $\sigma_+^q = \frac{\theta^2 + q^2(1-\theta)^2 + qR}{V + 2qR + q^2V}$, $\sigma^q = \frac{(R/2) + q(1-\theta)^2}{R + qV}$ and $\sigma_- = \frac{(1-\theta)^2}{V}$, where $R \equiv 2\theta(1-\theta)$ and $V \equiv \theta^2 + (1-\theta)^2$. It is easy to verify that rational agents discount the α signals more than the \emptyset signals because the α signals can be forged: $\sigma^q \leq \frac{1}{2} = \sigma^G$ and $\sigma_+^q \leq \sigma_+^G$. This effect becomes stronger as q grows: $\partial\sigma^q/\partial q < 0$ and $\partial\sigma_+^q/\partial q < 0$.

decision is state-dependent. This assumption is important because it ensures the government bears a cost for manipulating information.

Together, Assumptions 1 and 2 imply that the government and the public always agree when the evidence is clear cut. They both favor implementation when $(\alpha, \alpha)^G$ and they both favor the status quo when $(\emptyset, \emptyset)^G$. Nevertheless, disagreement can arise when the signals are not genuine, or when genuine signals are mixed. In the latter case, the public supports the status quo (by Assumption 1), but the government may prefer implementation (if B is sufficiently large). This disagreement is the source of the agency problem in our setting.

Observability of the Agency’s Bias. An important issue is whether the public can observe the agency’s bias q . For most of the paper, we will focus on the polar opposite scenarios of ‘transparency’ and ‘secrecy’. Under transparency, both the agency’s report and its levels of bias are observable, while under secrecy, neither is observable. Transparency should therefore be interpreted as an environment where information is easily accessible, not just about the contents of the report, but also about the staffing, track record and likely bias of the agency that drafts it. In contrast, under secrecy, information about the agency as well as the report is tightly guarded.⁸ Intermediate cases between transparency and secrecy are discussed in Section 6.

3. Transparency

We assume throughout this section that all information must be truthfully disclosed and that the agency is nonindependent. Since the government and the public both observe the signals \mathbf{s}^q and the agency’s bias q , they will share the same posterior beliefs. Let $\pi(p(\mathbf{s}^q), v(\mathbf{s}^q))$ denote the government’s payoff given \mathbf{s}^q , where $p(\mathbf{s}^q)$ denotes the policy decision of the government and $v(\mathbf{s}^q)$ denotes the decision supported by the public. Let $\Pr(\mathbf{s}^q|\mathbf{s}^G)$ be the probability of

⁸The focus on the polar opposite scenarios of transparency and secrecy can be justified on two additional grounds. First, news coverage following the publication of the report may put the agency in the public eye. Following disclosure, therefore, information about the likely bias of the agency may become newsworthy/observable.

Second, publication of the report may allow the reliability and bias of the agency’s sources to be assessed. A concrete example is provided by the British Government’s 2003 "dodgy dossier". This dossier claimed that it drew "upon a number of sources, including intelligence material". Soon after its publication, however, it was discovered by a Cambridge University lecturer to have plagiarized past academic articles to a large extent, including the work of an American research student. Although ministers were exonerated from the charge of having misled parliament (Foreign Affairs Committee, 2003), this dossier "undermined the credibility of [the Government’s] case for war" (Foreign Affairs Committee, 2003, p. 42) and arguably cast doubt on the competence and motivations of the civil servants in charge of drafting it.

observing \mathbf{s}^q conditional on genuine signals \mathbf{s}^G . For any given $q \in [0, 1]$, the government's expected payoff is

$$E(\pi^q) = \sum_{\mathbf{s}^G \in \{\alpha, \emptyset\}^2} \left[\sum_{\mathbf{s}^q \in \{\alpha, \emptyset\}^2} \pi(p(\mathbf{s}^q), v(\mathbf{s}^q) | \mathbf{s}^G) \Pr(\mathbf{s}^q | \mathbf{s}^G) \right] \Pr(\mathbf{s}^G). \quad (3.1)$$

Equation (3.1) shows that manipulating information affects the government through two distinct channels. Higher levels of q undermine the government's ability to tailor its policy decision $p(\mathbf{s}^q)$ to the true state. This harms the government because its optimal decision is state-dependent. However, changing the distribution of observed signals $\Pr(\mathbf{s}^q | \mathbf{s}^G)$ also allows the government to shape public opinion, $v(\mathbf{s}^q)$. This can benefit the government by helping convince the public to support implementation.

We begin with a preliminary result showing that the equilibrium level of bias is bounded from above.

Lemma 1. *The government will never choose a level of bias that leads the public to always support the status quo. Formally, in equilibrium, $q \in [0, q^{\max}]$, where $q^{\max} \in [0, 1)$ solves $\sigma_+^{q^{\max}} = C_a$.*

Because the public is rational, the weight it places on a positive report is decreasing in the level of bias. When $q > q^{\max}$, the bias is so large that the public disregards the report: citizens support the status quo even when both signals are positive. The government strictly prefers setting any $q \in [0, q^{\max}]$, which provides better information for decision making and can also generate support for implementation through a positive report.

Having restricted the set of q 's that can be optimal, we now examine which policy decisions are taken and supported in equilibrium. We begin with a partial result that simplifies the government's optimization problem. A full characterization of equilibrium play is provided later in Proposition 1.

Lemma 2. *In equilibrium, the government selects implementation when both signals are positive and the status quo when both signals are negative. The public supports implementation if and only if both signals are positive.*

Lemma 2 easily follows from Assumptions 1 and 2 and the fact that in equilibrium $q \leq q^{\max}$. It shows that the public and the government always agree on the appropriate course of action when the evidence is clear-cut (i.e., when the signals are both positive or both

negative). What Lemma 2 does not show is whether the government will implement the policy when the evidence is mixed, even though the public supports the status quo. To distinguish between the two relevant cases, we make the following definition.

Definition (discipline). Fix $q \leq q^{\max}$. The government is said to be *disciplined by public opinion* if it selects implementation if and only if both signals are positive.

A government that is disciplined by public opinion selects the status quo when the evidence is mixed, and so always enjoys public support.⁹ For given $q \leq q^{\max}$, let $E(\pi_d^q)$ denote the government's payoff under discipline. Specifically, let $E(\pi_d^q)$ be a special case of (3.1) where (i) $q \leq q^{\max}$, (ii) the government selects implementation if and only if both signals are positive, and (iii) the public supports implementation if and only if both signals are positive. Note that $E(\pi_d^q)$ incorporates all the requirements in Lemmas 1 and 2 as well as the notion of discipline.

A government that is not disciplined by public opinion will select implementation when the evidence is mixed, despite a lack of public support. Let $E(\pi_{nd}^q)$ be the government's payoff in that case. Thus, $E(\pi_{nd}^q)$ is a special case of (3.1), where (i) $q \leq q^{\max}$, (ii) the government selects implementation if and only if the signals are positive or mixed, and (iii) the public supports implementation if and only if both signals are positive. Both $E(\pi_d^q)$ and $E(\pi_{nd}^q)$ are explicitly computed in Appendix A.

To simplify the exposition of the results, we rule out corner solutions that would arise when the constraint $q \leq q^{\max}$ is binding:

$$q^* \equiv \arg \max_{q \in [0,1]} E(\pi_d^q) \leq q^{\max}. \quad (3.2)$$

Like Assumption 2, condition (3.2) requires that B not to be too large. An explicit condition is provided in the appendix (see the proof of Proposition 1).

We can now state this section's main result.

Proposition 1. *In the equilibrium of the transparency game, the public supports implementation if and only if $\mathbf{s}^q = (\alpha, \alpha)$. Moreover*

- i. *If $B \leq C_a - \frac{1}{2}$, then the government selects implementation if and only if $\mathbf{s}^q = (\alpha, \alpha)$, and the equilibrium level of bias is zero.*

⁹This is true because $q \leq q^{\max}$. If $q > q^{\max}$, then the public would never support implementation.

- ii. If $B \in (C_a - \frac{1}{2}, C_a - \frac{1}{2} + E)$, then the government selects implementation if and only if $\mathbf{s}^q = (\alpha, \alpha)$, and the equilibrium level of bias is $q^* = \frac{R}{V} \frac{B - C_a + 1/2}{C_a - B - \sigma_-} \leq q^{\max}$, where $R \equiv 2\theta(1 - \theta)$ and $V \equiv \theta^2 + (1 - \theta)^2$.
- iii. If $B \geq C_a - \frac{1}{2} + E$, then there are two possible cases. In the first case, the government selects implementation if and only if $\mathbf{s}^q \neq (\emptyset, \emptyset)$, and the level of bias is zero. In the second case, the government selects implementation if and only if $\mathbf{s}^q = (\alpha, \alpha)$, and the level of bias is $q^* = \frac{R}{V} \frac{B - C_a + 1/2}{C_a - B - \sigma_-} \leq q^{\max}$. The first case arises if and only if $E(\pi_d^{q^*}) \leq E(\pi_{nd}^G)$.

Proposition 1 fully characterizes equilibrium play in the transparency game. Case (i) deals with a situation where B is so small that, regardless of electoral concerns, the government would select the status quo when the evidence was mixed.¹⁰ The government's interests are aligned with those of the public, so there is no need to manipulate information.

Case (ii) deals with a situation where the government would select implementation if $E = 0$ and the genuine signals were mixed, ($B > C_a - \frac{1}{2}$), but where electoral concerns leave it unwilling to make an unpopular decision ($B < C_a - \frac{1}{2} + E$). Thus, there is a conflict of interest between the government and the public, but the government still selects the public's preferred policy because of electoral concerns.

However, precisely because public opinion is so powerful, the government now has an incentive to shape it. Note that the government's choice of q affects the distribution of the observed signals \mathbf{s}^q . This has two effects on the government's payoff $E(\pi_d^q)$. On the one hand, higher levels of bias q reduce the quality of information available for decision making. Specifically, with probability $\frac{1}{2}q^2V$, two genuine negative signals are transformed into two positive signals. This will result in the policy being implemented and in an expected loss for the government of $C_a - B - \sigma_-$, relative to the counterfactual where $q = 0$.¹¹ On the other hand, the government wants to 'trick' the public into supporting implementation when the genuine signals are mixed. Manipulations help the government because they can transform mixed signals into two positive signals. This occurs with probability qR and yields a net benefit of $B - C_a + \frac{1}{2}$ to the government, relative to the counterfactual where $q = 0$.¹² The

¹⁰This follows from equation (2.2) and the fact that $\sigma_{Gov} \leq \frac{1}{2}$ when the evidence is mixed.

¹¹In fact, $\Pr((\alpha, \alpha)^q | (\emptyset, \emptyset)^G) \Pr((\emptyset, \emptyset)^G) = \frac{1}{2}q^2V$ and $\pi(w, w | (\emptyset, \emptyset)^G) - \pi(n, n | (\emptyset, \emptyset)^G) = -[C_a - B - \sigma_-]$, where $-[C_a - B - \sigma_-] > 0$ by Assumption 2.

¹²In fact, $2 \Pr((\alpha, \alpha)^q | (\emptyset, \alpha)^G) \Pr((\emptyset, \alpha)^G) = qR$ and $\pi(w, w | (\emptyset, \alpha)^G) - \pi(n, n | (\emptyset, \alpha)^G) = B - C_a + \frac{1}{2}$.

optimum q balances precisely these gains from manipulation against the costs associated with poor decision making.¹³

The third case is when $B \geq C_a - \frac{1}{2} + E$. This case is more complicated because the government may or may not select implementation when the evidence is mixed. Condition $B \geq C_a - \frac{1}{2} + E$ implies that in the absence of bias, $q = 0$, the government would select implementation (no discipline). However, q is endogenous, and the government's optimal level of bias may differ from zero. The crucial observation here is that for any given signal realization \mathbf{s}^q , the government's incentive to select implementation is (weakly) decreasing in q . As q becomes large, observed α signals are more likely to be fake, leaving the government more reluctant to implement the policy. A large level of bias may therefore lead the government to prefer the status quo when the evidence is mixed (discipline).

We distinguish between two cases, depending on whether the optimal q is above or below a cutoff \hat{q} . The cutoff is defined so that, if $q \leq \hat{q}$, the government selects implementation after observing a mixed report (no discipline).¹⁴ Thus, the government payoff on $[0, \hat{q}]$ is $E(\pi_{nd}^q)$. Conversely, if $q \geq \hat{q}$, then the government selects the status quo after observing a mixed report (discipline). Thus, the government payoff on $q \in (\hat{q}, q^{\max}]$ is $E(\pi_d^q)$.

Taken together, we obtain the following expression for the government's payoff as a function of q when $B \geq C_a - \frac{1}{2} + E$:

$$E(\pi_{LB}^q) \equiv \begin{cases} E(\pi_{nd}^q) & \text{if } q \in [0, \hat{q}] \\ E(\pi_d^q) & \text{if } q \in (\hat{q}, q^{\max}] \end{cases}. \quad (3.3)$$

Thus, when private benefits B are large (LB), the government will maximize $E(\pi_{LB}^q)$ with respect to $q \in [0, q^{\max}]$. Two types of equilibria can arise: one characterized by a low bias in information and by no discipline (when the optimum q lies on the interval $[0, \hat{q}]$), and another characterized by a large bias in information and by discipline (when the optimum q lies on the interval $(\hat{q}, q^{\max}]$). In the proof of Proposition 1, we show that the government's payoff is convex over $[0, \hat{q}]$ and concave over $(\hat{q}, q^{\max}]$. It is because of this nonconcavity that, as stated in Proposition 1(iii), the optimum level of the bias is either zero or $q^* \equiv \arg \max_{q \in [0, 1]} E(\pi_d^q)$.

¹³Manipulations can also transform two negative signals into two mixed signals. Under discipline, however, this change is inconsequential because the government always selects the status quo.

¹⁴Formally, \hat{q} is defined as the minimum between 1 and the solution to $\sigma^q = C_a - B + E$. See Appendix B for an explicit characterization.

3.1. Can the Public Benefit from Biased Information?

Proposition 1 shows that, depending on parameter values, two different equilibrium outcomes can arise: one with no discipline and no bias in information ($q = 0$), and another with discipline and a positive bias in information ($q = q^*$). Moreover, under Case (iii), the government may be willing to choose discipline precisely because it is also able to manipulate information. If the bias were forced to be zero, so that the agency was independent, then the government would become less cautious and select implementation when the evidence was mixed.

This subsection explores the idea that, due to the positive equilibrium association between bias and discipline, the government's ability to manipulate information may benefit the public. We begin with an illustrative example, describing the region of parameter space (B, E) for which this is indeed the case.

The following figure shows the equilibrium level of discipline and bias implied by Proposition 1, as a function of B and E , when $C_a = 0.6$ and $\theta = 0.8$.




Figure 1: Equilibrium Level of Discipline and Bias, for $C_a = 0.6$ and $\theta = 0.8$

Figure 1 illustrates the regions of parameter space corresponding to the various cases of Proposition 1. Below the horizontal line $B = C_a - 1/2$, Case (i) implies the government will be disciplined by public opinion and will choose zero bias. Above this horizontal line but

below the 45 degree line $B = C_a - 1/2 + E$, Case (ii) implies the government will be disciplined by public opinion, choose positive bias q^* , and would remain disciplined even if it were forced to set a bias of zero. Above the 45 degree line are the two subregions corresponding to Case (iii): one where the government chooses no discipline and no bias, and another where it chooses discipline and positive bias q^* .¹⁵ It is in this latter subregion that the government's ability to manipulate information may benefit the public, which is the case in the shaded area.

Whithin this latter subregion of Case (iii), the government's ability to manipulate information has two effects on citizen welfare. Bias makes positive signals less reliable, so that the government selects the status quo when the evidence is mixed. This moderating effect of bias helps the public by making the government more cautious. However, bias also means that seemingly positive signals may be forgeries, which hurts the public by unduly stacking the deck in favour of implementation. The public will benefit from the government's ability to manipulate information if the gains from discipline generated by this moderating effect outweigh the losses due to biased decision making. Moreover, these losses are increasing in the level of bias. Because q^* is increasing in B and independent of E , there is a threshold value of B below which the public is willing to accept bias q^* to more closely align the government's interests with its own.¹⁶ The public benefits from the government's ability to manipulate information in the shaded area of Figure 1, which is the part of subregion (iii - Discipline and Bias) where q^* is sufficiently low.

More generally, as shown in Proposition 2 below, there is always a region of parameter space (B, E) corresponding to the shaded area in Figure 1, where the public strictly benefits from the government's ability to manipulate information.

Proposition 2. *For any C_a and $\theta \in (\frac{1}{2}, 1)$, there are values of B and E such that a commitment not to manipulate information strictly hurts the public. Specifically, there exists $\bar{B} > C_a - \frac{1}{2}$, and $\underline{E}(B) < B - C_a + \frac{1}{2}$ for any $B \in (C_a - \frac{1}{2}, \bar{B})$, such that $E(U^{q^*}) > E(U^G)$ if and only if $B \times E \in (C_a - \frac{1}{2}, \bar{B}) \times (\underline{E}(B), B - C_a + \frac{1}{2}]$.¹⁷*

¹⁵Figure 1 also shows that the government will choose discipline whenever E exceeds a certain lower bound. Intuitively, an increase in E makes discipline more attractive by increasing the impact of public opinion.

¹⁶The threshold value of B is identified by $E(U_d^{q^*}) = E(U_{nd}^G)$, where B only affects the public's payoff via $q^* = \frac{R}{V} \frac{B - C_a + 1/2}{C_a - B - \sigma_-}$.

¹⁷Proposition 2 takes into account condition (3.2), $q^* \leq q^{\max}$. See the appendix for more details.

Finally, note that our analysis shows that bias can only moderate government policy if electoral concerns are relatively weak. When $E > B - C_a + \frac{1}{2}$, the government is always disciplined by public opinion, regardless of the level of bias. Electoral concerns then suffice to ensure that the government caters to the public. Thus, from the public's point of view, manipulations simply stack the deck in favor of implementation, which decreases their payoff. The present model therefore suggests that the independence of government agencies should unambiguously benefit the public in mature democracies, where E is large. In contrast, in less mature democracies, where governments care about public opinion but are not fully responsive to electoral concerns, non independence may sometimes be socially optimal.

4. Secrecy

The previous section studied the case where the government must truthfully disclose both the signal realizations and the level of bias. This section analyzes the polar opposite scenario of *secrecy*: the government commits not to disclose either \mathbf{s}^g or q . The main complication that arises is that, as in Canes et al. (2001) and Maskin and Tirole (2004), the government now has private information. As a result, the public's choice will in general depend on the policy decision of the government, which potentially conveys information.¹⁸

Despite this complication, we can characterize equilibrium play. Let $R = 2\theta(1 - \theta)$ and $V = \theta^2 + (1 - \theta)^2$. Furthermore, let $\hat{\sigma} = \frac{\theta^2 + R}{1 + R} \in (\frac{1}{2}, 1)$ be the public's belief that A is the true state when $q = 0$, implementation is selected by the government, and the government is disciplined by public opinion.

Proposition 3. *The following is a perfect Bayesian equilibrium of the secrecy game. The level of bias is zero. The government selects implementation when both signals are positive and selects the status quo when both signals are negative. The public supports the status quo whenever the status quo is selected.*

- i. *If $B \leq C_a - \frac{1}{2}$, then the government selects the status quo when the evidence is mixed and the public supports implementation whenever implementation is selected.*

¹⁸On the other hand, the public's choice v cannot depend on the realization of the signals \mathbf{s}^g or the bias q because they are unobservable.

- ii. If $B \in (C_a - \frac{1}{2}, C_a - \frac{1}{2} + E)$ and $\hat{\sigma} < C_a$, then the government selects implementation with probability $\check{s} = \frac{V}{R} \frac{\sigma_+^G - C_a}{2(C_a - 1/2)}$ when the evidence is mixed and the public supports implementation with probability $1 - \frac{1/2 - (C_a - B)}{E}$ when implementation is selected.
- iii. If $B \in (C_a - \frac{1}{2}, C_a - \frac{1}{2} + E)$ and $\hat{\sigma} \geq C_a$, then the government selects implementation when the evidence is mixed and the public supports implementation whenever implementation is selected.
- iv. If $B \geq C_a - \frac{1}{2} + E$ and $\hat{\sigma} < C_a$, then the government selects implementation when the evidence is mixed and the public supports the status quo whenever implementation is selected.
- v. If $B \geq C_a - \frac{1}{2} + E$ and $\hat{\sigma} \geq C_a$, then the government selects implementation when the evidence is mixed and the public supports implementation whenever implementation is selected.

The equilibrium in Proposition 3 exhibits several intuitive features. First, consistent with conventional wisdom, secrecy is shown to be effective at protecting the integrity of the decision-making process. The government has no incentive to set a positive bias because neither q nor the signal realizations are observed by the public. Increasing q simply reduces the quality of information available to the government, so the equilibrium level of bias is zero.

A second intuitive feature of the equilibrium is that, as the government's private benefits B grows large, the government is less likely to be disciplined by public opinion. Proposition 3 shows that the government always selects implementation when the signals are positive, and always selects the status quo when the signals are negative. Thus, for the government to be disciplined by public opinion, we only need to check whether the government selects the status quo when the signals are mixed. Proposition 3 shows that, when the signals are mixed, implementation is always selected when B is large (cases (iv)-(v)), and it is often selected when B is intermediate (cases (ii)-(iii)). It is only when B is small that the status quo is always selected (case (i)). Thus, as B grows large, the government is less likely to cater to public opinion.

It is also instructive to compare the equilibrium outcomes under transparency and secrecy (Propositions 1 and 3). More cases must be distinguished under secrecy than under transparency (five versus three). Under secrecy, when the government selects implementation,

the public cannot observe whether the signal are positive or mixed. The condition $\hat{\sigma} < C_a$ describes scenarios where mistaken implementation is sufficiently costly for the public to then support the status quo. When mistaken implementation is less costly ($\hat{\sigma} \geq C_a$), the situation is reversed. These complications do not arise under transparency because the public can observe the signals, so that beliefs $\hat{\sigma}$ play no role.

Note also that, because public opinion is influenced by the government's policy decision, a mixed strategy equilibrium can arise under secrecy. This does not happen under transparency.¹⁹

Finally, from the public's point of view, the choice between transparency and secrecy involves a key trade-off between manipulations and discipline. Manipulations are always (weakly) lower under secrecy, while discipline is always (weakly) higher under transparency. That manipulations are lower under secrecy is obvious as $q = 0$. Let us therefore compare transparency and secrecy in terms of discipline. When there is no conflict of interest (case (i)), the government is disciplined by public opinion under both scenarios. In contrast, when the conflict of interest is intermediate, the government is always disciplined by public opinion under transparency (Proposition 1(ii)) but not under secrecy (Proposition 3(ii)-(iii)). Furthermore, when the conflict of interest is large, the government is sometimes disciplined by public opinion under transparency (Proposition 1(iii)) but never under secrecy (Proposition 3(iv)-(v)). Thus discipline is always at least as likely under transparency as under secrecy.

This lack of discipline under secrecy is caused by a relative lack of accountability. Without observing the report, citizens cannot determine exactly why a particular decision was taken. For example, the government's decision to select implementation could be based on strong

¹⁹To see how mixed strategies can arise in case (ii), recall that when $B \in (C_a - \frac{1}{2}, C_a - \frac{1}{2} + E)$, the government would like to select implementation when the evidence is mixed (since $B > C_a - \frac{1}{2}$), but because E is large relative to B ($C_a - \frac{1}{2} + E > B$), it is unwilling to select an unpopular policy. Now, given that the report is not disclosed, can the government get away with selecting implementation when the evidence is mixed? (When the report is disclosed, the answer is no because the public supports the status quo if the evidence is mixed. Thus discipline obtains. See Proposition 1(ii).) Suppose the government does select implementation when the evidence is mixed. Then after observing implementation, the public must believe that the probability that the true state is A is $\hat{\sigma}$. If $\hat{\sigma} \geq C_a$, this probability is high enough to induce the public to support implementation. Thus the government gets away with selecting implementation and the equilibrium is in pure strategies (case (iii)). However, if $\hat{\sigma} < C_a$, then the public does not support implementation. Because $B < C_a - \frac{1}{2} + E$, selecting implementation with probability one is not optimal for the government. Selecting the status quo with probability one is also not optimal. Note in fact that if implementation is only selected when the signals are positive, then the public must always support the government. This leads to a contradiction because if the government policy is always supported, then the government follows its intrinsic preferences and selects implementation with probability one when the evidence is mixed. Thus, when $\hat{\sigma} < C_a$, the only equilibrium can be in mixed strategies (case (ii)).

evidence ($\mathbf{s}^G = (\alpha, \alpha)$) or mixed evidence ($\mathbf{s}^G = (\alpha, \emptyset)$). The public would only like to punish the government in the latter case but it cannot do so without seeing the report. As a result, the government is less accountable and thus less responsive to public desires. It is easy to construct examples where, because of this trade-off between manipulations and accountability, either transparency or secrecy is preferred by the public.

5. Independence and Optimal Constitutions

So far we have assumed that the government can easily interfere with the workings of the agency in charge of collecting information. This is a reasonable assumption if, as in the U.S., the President appoints and can remove the heads of the executive agencies, thus exerting enormous influence over their policy decisions. Sometimes, however, executive influence over government agencies is more limited. Of special interest is the case of *independent* agencies such as the Federal Trade Commission in the U.S. and the Bank of England in the U.K. These agencies are not subject to the same degree of political control as other executive agencies and are insulated from political pressure, for instance by limiting the removal of their heads to certain causes.

This section considers the implications of granting full independence to the government agencies in charge of collecting information. Formally, independence is modelled as a commitment not to manipulate information. Thus, an independent agency will carry out its job as objectively as possible. We first compare transparency and secrecy under the assumption that the agency is independent.

Proposition 4. *Suppose the government agency is independent (i.e., $q \equiv 0$). Then the public's payoff is always higher under transparency than under secrecy.*

The intuition for this result is straightforward: if information cannot be manipulated, then only accountability matters, and transparency is in the interests of the public.

Having established this benchmark result, we now consider the more interesting case where both the disclosure rule (transparency or secrecy) and the degree of insulation of the government agency (independence or nonindependence) can be chosen to maximize the public's welfare. Following previous work, we refer to the stage when society decides the rules of the game as the 'constitutional' stage. The four constitutions we consider are shown in Table I.

Table I: Constitutions

I	Transparency & Independent Agency	II	Transparency & Nonindependent Agency
III	Secrecy & Independent Agency	IV	Secrecy & Nonindependent Agency

A constitution is said to be optimal if it maximizes the public’s welfare. Proposition 5 characterizes optimal constitutions.

Proposition 5. *An optimal constitution always involves transparency. The comparison between Constitution I (transparency & independent agency) and Constitution II (transparency & nonindependent agency) is ambiguous.*

In an environment where information disclosure creates incentives for manipulation, it is perhaps surprising that transparency is always optimal. The intuition for this result is simple: the chief advantage of secrecy – unbiased information – can more effectively be achieved by insulating the government agency from political pressure. To see this more formally, note that by Proposition 4, Constitution I (transparency & independent agency) dominates Constitution III (secrecy & independent agency). Moreover, the two constitutions involving secrecy (Constitutions III and IV) are payoff equivalent because under secrecy q is always equal to zero in equilibrium. Thus transparency (Constitution I) always dominates secrecy (Constitutions III and IV).²⁰

However, this result does not imply that granting independence to government agencies is necessarily in the public interest. Biased information can induce the government to behave more cautiously, thus mitigating the agency conflict between the government and the public (see Propositions 1 and 2). As a result, the comparison between Constitution I and Constitution II is ambiguous.

6. Extensions

The analysis so far has focused on the polar opposite scenarios of transparency and secrecy. Under transparency both the report and the bias in information are observable, while under

²⁰This result requires the combination of transparency and independence to be available at the constitutional stage. If independence was not feasible, the trade-off between manipulations and accountability highlighted in the previous section would obviously reappear.

secrecy neither is observable. This section briefly discusses two intermediate scenarios, disclosure of the report with unobservable bias and nondisclosure of the report with observable bias, as well as a third scenario where disclosure is voluntary. More details about the results in this section are available from the authors upon request.

Disclosure with Unobservable Bias. We begin with the case when the report is disclosed but the public does not observe the bias of the agency. It can be shown that if assumption (3.2) holds, then the equilibrium outcome when q is unobservable is exactly the same as when q is observable. Thus, under assumption (3.2), Proposition 1 is not affected by the unobservability of q .

The idea behind this result is simple. When q is not observable, the public must form some conjecture about the level of bias chosen by the government. In equilibrium, this conjecture must be correct. Assumption (3.2) ensures that, when the public believes that $q \leq q^{\max}$, then the government will actually choose a level of bias $q \leq q^{\max}$. Thus the public's belief can be made consistent with the play of the game. In particular, the government's incentives to set any $q < q^{\max}$ are just as in Section 3, so the same equilibrium as in Proposition 1 (with observable q) can be supported.²¹

Nondisclosure with Observable Bias. An alternative scenario arises when the government commits not to disclose the report but the agency bias is observable. This scenario is plausible if the public is well-informed about the reputation and policy dispositions of the individuals working for the agency, even though the report is not disclosed.

Compared to the case where bias is unobservable (secrecy), an interesting new effect can arise. Specifically, the government may choose a strictly positive bias to commit itself to a more congruent decision rule. Intuitively, by appointing a head of the agency who is well-known to be biased in favor of implementation, the government can credibly commit not to select implementation when the undisclosed signals are mixed. A mixed report from a very biased bureaucrat provides very little evidence in support of implementation.

The optimal choice of q is therefore determined by two conflicting effects. On the one hand, manipulations reduce the quality of information available to the government, which reduces its payoff. On the other hand, a sufficiently high level of bias allows the government

²¹If assumption (3.2) does not hold, then equilibrium behavior is more complicated and will generally involve randomization over q .

to commit to a decision rule that the public prefers. This can induce the public to support the government policy more often, thereby increasing the government payoff.

That being said, this new effect of nondisclosure with observable bias is only present if the government needs to convince the public to support its decision. Proposition 3 shows that the public often supports the government's decision even if the agency is unbiased ($q = 0$) (see cases (i), (iii) and (v)). In this sense, these elements of Proposition 3 will continue to hold whether or not q is observable.

Voluntary Disclosure. We have assumed so far that the government must either disclose the contents of the report or must keep it secret. More commonly, however, policymakers have discretion as to whether to release information. We now consider a variant of the model where the government cannot commit to any disclosure rule: disclosure is voluntary. We argue that voluntary disclosure will effectively result in all information being disclosed. Indeed, since information is hard in this model, Milgrom's (1981) 'unraveling' result applies.

To see the logic of this result, suppose the public expects the government to disclose favorable information (the α signals). Thus nondisclosure is interpreted as evidence that the information is unfavorable (a \emptyset signal), which provides the government with a strong incentive to disclose favorable information. Specifically, a government that receives two positive signals will disclose them and implement the policy with public support. When the evidence is mixed or unfavorable, whether or not the signals are disclosed is inconsequential because the public would realize that at least one of them is unfavorable. Thus they would not support implementation. All information is therefore revealed, and the analysis would proceed as in Section 3.

7. Conclusion

This paper develops a model where disclosure of information gives the government an incentive to "fix the evidence" around its ex ante favored policy. Decision-relevant information is collected by an agency, but the government can distort this process, for instance by staffing the agency with biased individuals. The key trade-off the government faces is between protecting the quality of the information available for public decision making (if the agency is unbiased) and molding public opinion (if the agency is biased). Surprisingly, we find that insulating the agency from political pressure, so that the agency is always unbiased, is not

necessarily in the public interest. A biased information gathering process can in fact induce the government to act more cautiously in response to information supporting its ex ante preferred policy. This moderating effect of bias can more than outweigh the welfare losses caused by biased information.

We are not the first to study whether government agencies should be insulated from external or political pressures. Moe (1989, 1990) argues that government agencies are sometimes intentionally created to be unresponsive to political pressures to alleviate the risk of political power fluctuations. Prendergast (2003) points out that bureaucrats' tendency to inefficiently accede to customer demands may require appropriate organizational responses, such as insulating government agencies from customer complaints. Betts (2004) notes that a close connection between the President and top intelligence officials may be preferable to the lack of such a connection because the risks of insulation and unresponsiveness often far outweigh those of politicization. This paper highlights a novel drawback of bureaucratic independence: the risk that candid advice from government agencies may make policymakers very responsive to information supporting their ex ante favored policy, thus exacerbating the conflict of interest between the government and the public.

Acknowledgements: We would like to thank Diego Gambetta, Clare Leaver, Ian Jewitt, Meg Meyer, David Myatt, Torsten Persson, Andrea Prat, John Quah and numerous seminar participants for helpful discussions and remarks. Financial support from the British Academy is also gratefully acknowledged. All remaining errors are our own.

References

- Alesina, A., Tabellini, G. (2007). Bureaucrats or politicians? Part I: a single policy task. *American Economic Review*, vol. 97, pp. 169-179.
- Alesina, A., Tabellini, G. (2008). Bureaucrats or politicians? Part II: multiple policy tasks. *Journal of Public Economics*, vol. 92, pp. 426-447.
- Ashworth, S., Shotts, K.W. (2010). Does informative media commentary reduce politicians' incentives to pander? *Journal of Public Economics*, vol. 94, pp. 838-847.
- Banisar, D. (2004). 'The Freedominfo.org Global Survey. Freedom of Information and Access to Government Record Laws around the World', available at http://www.freedominfo.org/documents/global_survey2004.pdf

Besley, T. and Prat A. (2006). ‘Handcuffs for the Grabbing Hand? Media Capture and Political Accountability’, *American Economic Review*, vol. 96, pp. 720–736.

Betts, R.K. (2004). ‘The New Politics of Intelligence’, *Foreign Affairs*, vol. 83, pp. 2-8.

Biglaiser, G. and Mezzetti, C. (1997). ‘Politicians’ decision making with re-election concerns’, *Journal of Public Economics*, vol. 66, pp. 425-447.

Canes-Wrone, B., Herron M.C. and Shotts K.W. (2001). ‘Leadership and Pandering: A Theory of Executive Policymaking’, *American Journal of Political Science*, vol. 45, pp. 532-550.

Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction (2005) ‘Report to the President of the United States’, downloadable at <http://govinfo.library.unt.edu/wmd/report/index.html>

Committee of Privy Counsellors. Chairman: The Rt Hon The Lord Butler of Brockwell KG GCB CVO (2004). ‘Review of Intelligence on Weapons of Mass Destruction’, downloadable at <http://www.butlerreview.org.uk/>.

Cremer, J. (1995). ‘Arm’s Length Relationships’, *Quarterly Journal of Economics*, vol. 110, pp. 275-295.

Foreign Affairs Committee (2003). ‘The Decision to go to War in Iraq’, downloadable at <http://www.publications.parliament.uk/pa/cm200203/cmselect/cmfa/813/813.pdf>

Holmstrom, B. (1979). ‘Moral Hazard and Observability’, *Bell Journal of Economics*, vol.10, pp. 74-91.

Levy, G. (2007). ‘Decision Making in Committees: Transparency, Reputation and Voting Rules’, *American Economic Review*, vol. 97, pp. 150-168.

Marschak, J. and Radner R. (1972). *Economic Theory of Teams*, Yale University Press, New Haven and London.

Maskin, E. and Tirole J. (2004). ‘The Politician and the Judge’, *American Economic Review*, vol. 94, pp. 1034-1051.

Milbourn, T.T., Shockley, R.L. and Thakor, A.V. (2001). ‘Managerial Career Concerns and Investments in Information’, *Rand Journal of Economics*, vol. 32, pp. 334-351.

- Milgrom, P. (1981). 'Good News and Bad News: Representation Theorems and Applications', *Bell Journal of Economics*, vol. 12, pp. 380-391.
- Moe T.M. (1989). 'The Politics of Bureaucratic Structure', in John E. Chubb and Paul E. Peterson, eds., *Can the Government Govern?*. Washington, D.C.: The Brookings Institution.
- Moe T.M. (1990). 'Political Institutions: The Neglected Side of the Story', *Journal of Law, Economics, & Organization*, vol. 6, pp. 213-253.
- Prat, A. (2005). 'The Wrong Kind of Transparency', *American Economic Review*, vol. 104, pp. 946-59.
- Prat, A. (2006). The More Closely we are Watched, the Better we Behave?. In *Transparency: The Key to Better Governance?* (eds. Christopher Hood and David Heald). Oxford University Press.
- Prendergast, C. (1993). 'A Theory of "Yes Men"', *American Economic Review*, vol. 83, pp. 757-770.
- Prendergast, C. (2003). 'The Limits of Bureaucratic Efficiency', *Journal of Political Economy*, vol. 111, pp. 929-958.
- Roberts, A. (2006). 'Dashed Expectations: Governmental Adaptation to Transparency Rules' in *Transparency: The Key to Better Governance?*, C. Hood and D. Heald eds., Oxford University Press, Oxford.
- Select Committee on Intelligence United States Senate (2004). 'Report of the Select Committee on Intelligence on the U.S. Intelligence Community's Prewar Intelligence Assessments on Iraq', downloadable at <http://www.gpoaccess.gov/serialset/creports/iraq.html>
- Stasavage, D. (2004). 'Open-Door or Closed-Door? Transparency in Domestic and International Bargaining', *International Organization*, vol. 58, pp. 667-703.
- Suurmond, G., Swank O. and Visser B. (2004). 'On the Bad Reputation of Reputational Concerns', *Journal of Public Economics*, vol. 88, pp. 2817-2838.
- Swank, J., Swank, O. and Visser, B. (2008). 'How Committees of Experts Deal with the Outside World: Some Theory, and Evidence from the FOMC', *Journal of the European*

Economic Association, vol. 6, pp. 478-486.

Appendix A: Payoffs

This appendix derives explicit expressions for the government payoff and citizen welfare in the transparency case.

The Government Payoff. The government payoff is given by

$$E(\pi^q) = \sum_{\mathbf{s}^G \in \{\alpha, \emptyset\}^2} \left[\sum_{\mathbf{s}^q \in \{\alpha, \emptyset\}^2} \pi(p(\mathbf{s}^q), v(\mathbf{s}^q) | \mathbf{s}^G) \Pr(\mathbf{s}^q | \mathbf{s}^G) \right] \Pr(\mathbf{s}^G), \quad (\text{A1})$$

where $\pi(p(\mathbf{s}^q), v(\mathbf{s}^q) | \mathbf{s}^G)$ is the government expected payoff when the observed signals are \mathbf{s}^q and the genuine signals are \mathbf{s}^G , and q is the probability that a genuine \emptyset signal is transformed into fake α signal.

To compute the probabilities in (A1), note that

$$\begin{aligned} \Pr((\alpha, \alpha)^q) &= \Pr((\alpha, \alpha)^q | (\alpha, \alpha)^G) \Pr((\alpha, \alpha)^G) + \\ &\quad 2 \Pr((\alpha, \alpha)^q | (\alpha, \emptyset)^G) \Pr((\alpha, \emptyset)^G) + \Pr((\alpha, \alpha)^q | (\emptyset, \emptyset)^G) \Pr((\emptyset, \emptyset)^G), \end{aligned}$$

$$\Pr((\alpha, \emptyset)^q) = \Pr((\alpha, \emptyset)^q | (\alpha, \emptyset)^G) \Pr((\alpha, \emptyset)^G) + \Pr((\alpha, \emptyset)^q | (\emptyset, \emptyset)^G) \Pr((\emptyset, \emptyset)^G),$$

$$\Pr((\emptyset, \emptyset)^q) = \Pr((\emptyset, \emptyset)^q | (\emptyset, \emptyset)^G) \Pr((\emptyset, \emptyset)^G).$$

Moreover

$$\begin{aligned} \Pr((\alpha, \alpha)^q | (\alpha, \alpha)^G) &= 1, \quad \Pr((\alpha, \alpha)^q | (\alpha, \emptyset)^G) = q, \quad \Pr((\alpha, \alpha)^q | (\emptyset, \emptyset)^G) = q^2, \\ \Pr((\alpha, \emptyset)^q | (\alpha, \emptyset)^G) &= (1 - q), \quad \Pr((\alpha, \emptyset)^q | (\emptyset, \emptyset)^G) = q(1 - q), \\ \Pr((\emptyset, \emptyset)^q | (\emptyset, \emptyset)^G) &= (1 - q)^2, \end{aligned}$$

and

$$\begin{aligned} \Pr((\alpha, \alpha)^G) &= \Pr((\emptyset, \emptyset)^G) = \frac{1}{2} (\theta^2 + (1 - \theta)^2), \\ \Pr((\alpha, \emptyset)^G) &= \Pr((\emptyset, \alpha)^G) = \theta(1 - \theta). \end{aligned}$$

Because in equilibrium Lemmas 1 and 2 must hold, we compute $E(\pi^q)$ under the assumption that (i) the public supports implementation if and only if both signals are positive, (ii) the government selects implementation when both signals are positive, and (iii) the government selects the status quo when both signals are negative. Thus, in equilibrium only two cases

can emerge: either the government is disciplined by public opinion (thus it selects the status quo when the evidence is mixed) or the government is not disciplined by public opinion (thus it selects implementation when the evidence is mixed).

If the government is disciplined by public opinion, then (A1) is given by

$$\begin{aligned}
E(\pi_d^q) &= \frac{1}{2}V\pi(a, a|(\alpha, \alpha)^G) + qR\pi(a, a|(\alpha, \emptyset)^G) + \frac{1}{2}q^2V\pi(a, a|(\emptyset, \emptyset)^G) \\
&+ (1 - q)R\pi(n, n|(\alpha, \emptyset)^G) + q(1 - q)V\pi(n, n|(\emptyset, \emptyset)^G) \\
&+ \frac{1}{2}(1 - q)^2V\pi(n, n|(\emptyset, \emptyset)^G),
\end{aligned} \tag{A2}$$

where

$$\begin{aligned}
V &\equiv \theta^2 + (1 - \theta)^2 = \Pr((\alpha, \alpha)^G) + \Pr((\emptyset, \emptyset)^G), \\
R &\equiv 2\theta(1 - \theta) = 2\Pr((\alpha, \emptyset)^G),
\end{aligned}$$

(the subscript d stands for discipline). It is simple to see, in fact, that $\Pr((\alpha, \alpha)^q | (\alpha, \alpha)^G) \Pr((\alpha, \alpha)^G) = \frac{1}{2}V$, $2\Pr((\alpha, \alpha)^q | (\alpha, \emptyset)^G) \Pr((\alpha, \emptyset)^G) = qR$, and so forth.

By contrast, if the government is not disciplined by public opinion, then (A1) is given by

$$\begin{aligned}
E(\pi_{nd}^q) &= \frac{1}{2}V\pi(a, a|(\alpha, \alpha)^G) + qR\pi(a, a|(\alpha, \emptyset)^G) + \frac{1}{2}q^2V\pi(a, a|(\emptyset, \emptyset)^G) \\
&+ (1 - q)R\pi(a, n|(\alpha, \emptyset)^G) + q(1 - q)V\pi(a, n|(\emptyset, \emptyset)^G) \\
&+ \frac{1}{2}(1 - q)^2V\pi(n, n|(\emptyset, \emptyset)^G),
\end{aligned} \tag{A3}$$

(the subscript nd stands for no discipline). Equations (A2) and (A3) are of course identical except when mixed signals are observed.

Computing the conditional payoffs π is straightforward. For instance, $\pi(a, a|(\alpha, \alpha)^G) = -C_a(1 - \sigma_+^G) + B = -C_a(1 - \frac{\theta^2}{V}) + B$, $\pi(a, n|(\alpha, \emptyset)^G) = -C_a(1 - \sigma^G) + B - E = -\frac{1}{2}C_a + B - E$, $\pi(n, n|(\emptyset, \emptyset)^G) = -C_n\sigma_- = -\frac{(1-\theta)^2}{V}C_n$ and so on. Plugging these values into (A2) and (A3) yields

$$\begin{aligned}
E(\pi_d^q) &= \frac{1}{2}V \left[-C_a \left(1 - \frac{\theta^2}{V} \right) + B \right] + qR \left[-\frac{1}{2}C_a + B \right] + \frac{1}{2}q^2V \left[-C_a \left(1 - \frac{(1-\theta)^2}{V} \right) + B \right] \\
&+ (1 - q)R \left[-\frac{1}{2}C_n \right] + q(1 - q)V \left[-\frac{(1-\theta)^2}{V}C_n \right] \\
&+ \frac{1}{2}(1 - q)^2V \left[-\frac{(1-\theta)^2}{V}C_n \right],
\end{aligned} \tag{A4}$$

and

$$\begin{aligned}
E(\pi_{nd}^q) &= \frac{1}{2}V \left[-C_a \left(1 - \frac{\theta^2}{V} \right) + B \right] + qR \left[-\frac{1}{2}C_a + B \right] + \frac{1}{2}q^2V \left[-C_a \left(1 - \frac{(1-\theta)^2}{V} \right) + B \right] \\
&\quad + (1-q)R \left[-\frac{1}{2}C_a + B - E \right] + q(1-q)V \left[-C_a \left(1 - \frac{(1-\theta)^2}{V} \right) + B - E \right] \quad (\text{A5}) \\
&\quad + \frac{1}{2}(1-q)^2V \left[-\frac{(1-\theta)^2}{V}C_n \right].
\end{aligned}$$

To simplify the computations in Appendix B, it is helpful to normalize $E(\pi_d^q)$ and $E(\pi_{nd}^q)$ by subtracting $E(\pi_{nd}^G)$ from both. Since

$$E(\pi_{nd}^G) = \frac{1}{2}V \left[-C_a \left(1 - \frac{\theta^2}{V} \right) + B \right] + R \left[-\frac{1}{2}C_a + B - E \right] + \frac{1}{2}V \left[-\frac{(1-\theta)^2}{V}C_n \right], \quad (\text{A6})$$

(simply set $q = 0$ in (A5)), we obtain

$$\Delta\pi_{d,nd}^q = qRE - \frac{1}{2}q^2V [C_a - B - \sigma_-] + (1-q)R \left[C_a - B + E - \frac{1}{2} \right], \quad (\text{A7})$$

$$\Delta\pi_{nd,nd}^q = qRE - \frac{1}{2}q^2V [C_a - B - \sigma_-] - (1-q)qV [C_a - B + E - \sigma_-], \quad (\text{A8})$$

where $\sigma_- = \frac{(1-\theta)^2}{V}$.

Citizen Welfare. Next, we derive the public's payoff (citizen welfare) in the transparency case. Citizen welfare is given by

$$E(U^q) = \sum_{\mathbf{s}^G \in \{\alpha, \emptyset\}^2} \left[\sum_{\mathbf{s}^q \in \{\alpha, \emptyset\}^2} U(p(\mathbf{s}^q), v(\mathbf{s}^q) | \mathbf{s}^G) \Pr(\mathbf{s}^q | \mathbf{s}^G) \right] \Pr(\mathbf{s}^G), \quad (\text{A9})$$

where $U(p(\mathbf{s}^q), v(\mathbf{s}^q) | \mathbf{s}^G)$ denotes the public's payoff when the observed signals are \mathbf{s}^q and the genuine signals are \mathbf{s}^G .

Because in equilibrium Lemmas 1 and 2 must hold, we also compute $E(U^q)$ under the assumption that (i) the public supports implementation if and only if both signals are positive, (ii) the government selects implementation when both signals are positive, and (iii) the government selects the status quo when both signals are negative. Again, two cases can arise.

If the government is disciplined by public opinion, then (A9) becomes

$$\begin{aligned}
E(U_d^q) &= \frac{1}{2}V \left[-C_a \left(1 - \frac{\theta^2}{V} \right) \right] + qR \left[-\frac{1}{2}C_a \right] + \frac{1}{2}q^2V \left[-C_a \left(1 - \frac{(1-\theta)^2}{V} \right) \right] \quad (\text{A10}) \\
&\quad + (1-q)R \left[-\frac{1}{2}C_n \right] + q(1-q)V \left[-\frac{(1-\theta)^2}{V}C_n \right] \\
&\quad + \frac{1}{2}(1-q)^2V \left[-\frac{(1-\theta)^2}{V}C_n \right],
\end{aligned}$$

(as always, the subscript d stands for discipline).²²

If the government is disciplined by public opinion, then (A9) becomes

$$\begin{aligned} E(U_{nd}^q) &= \frac{1}{2}V \left[-C_a \left(1 - \frac{\theta^2}{V} \right) \right] + qR \left[-\frac{1}{2}C_a \right] + \frac{1}{2}q^2V \left[-C_a \left(1 - \frac{(1-\theta)^2}{V} \right) \right] \\ &+ (1-q)R \left[-\frac{1}{2}C_a \right] + q(1-q)V \left[-C_a \left(1 - \frac{(1-\theta)^2}{V} \right) \right] \\ &+ \frac{1}{2}(1-q)^2V \left[-\frac{(1-\theta)^2}{V}C_n \right], \end{aligned} \quad (\text{A11})$$

(the subscript nd stands for no discipline).

We also normalize (A10) and (A11) by subtracting $E(U_{nd}^G)$ from both. This yields

$$\Delta U_{nd,nd}^q = -(1-q)qV(C_a - \sigma_-) - \frac{1}{2}q^2V(C_a - \sigma_-), \quad (\text{A12})$$

$$\Delta U_{d,nd}^q = (1-q)R \left(C_a - \frac{1}{2} \right) - \frac{1}{2}q^2V(C_a - \sigma_-). \quad (\text{A13})$$

Appendix B: Proofs

Proof of Lemma 1. To prove Lemma 1, it suffices to show that setting any $q > q^{\max}$ is dominated by setting $q = 0$. The government payoff when $q = 0$ is the maximum between $E(\pi_d^G)$ and $E(\pi_{nd}^G)$. To derive the government payoff when some $q > q^{\max}$ is selected, we use two facts. First, if $q > q^{\max}$, then, by definition of q^{\max} , the public never supports implementation. Second, if $\mathbf{s}^q = (\emptyset, \emptyset)$, then by Assumption 2 the government must choose the status quo. Thus, when $q > q^{\max}$, only three cases must be considered.

Case (i): the government selects implementation if and only if $\mathbf{s}^q = \{(\alpha, \alpha), (\alpha, \emptyset), (\emptyset, \alpha)\}$ (no discipline). Let $E(\pi_{nd,ns}^q)$ denote the government payoff in this case (the subscript ns is used to emphasize that when $q > q^{\max}$ the public never supports implementation). Note that because the government selects the same policies as in the no discipline case, $E(\pi_{nd,ns}^q)$ is equal to $E(\pi_{nd}^q)$, except that now the public does not support implementation when the signals are both positive. Thus

$$E(\pi_{nd,ns}^q) = E(\pi_{nd}^q) - \left(\frac{1}{2}V + qR + \frac{1}{2}q^2V \right) E,$$

²²The probabilities in (A9) have been computed above when deriving the government payoff. The payoffs conditional on the true underlying signals are also easy to derive. For instance, $U(a, a | (\alpha, \alpha)^G) = -C_a(1 - \sigma_+^G) = -C_a(1 - \frac{\theta^2}{V})$, $U(a, n | (\alpha, \emptyset)^G) = -C_a(1 - \sigma^G) = -\frac{1}{2}C_a$, $U(n, n | (\emptyset, \emptyset)^G) = -C_n\sigma_- = -\frac{(1-\theta)^2}{V}C_n$, and so forth.

since $\Pr((\alpha, \alpha)^q) = \frac{1}{2}V + qR + \frac{1}{2}q^2V$. Simple algebra yields

$$E(\pi_{nd,ns}^q) - E(\pi_{nd}^G) = -\frac{1}{2}VE - \frac{1}{2}q^2V [C_a - B + E - \sigma_-] - (1 - q)qV [C_a - B + E - \sigma_-] < 0.$$

Thus, if case (i) applies, then setting any $q > q^{\max}$ is strictly dominated by setting $q = 0$.

Case (ii): the government selects implementation if and only if $\mathbf{s}^q = (\alpha, \alpha)$ (discipline). Let $E(\pi_{d,ns}^q)$ denote the government payoff in this case. $E(\pi_{d,ns}^q)$ is equal to $E(\pi_d^q)$ except that now the public does not support implementation when the signals are both positive. Thus

$$E(\pi_{d,ns}^q) = E(\pi_d^q) - \left(\frac{1}{2}V + qR + \frac{1}{2}q^2V \right) E.$$

Simple algebra yields

$$E(\pi_{d,ns}^q) - E(\pi_{nd}^G) = -\frac{1}{2}VE - \frac{1}{2}q^2V [C_a - B + E - \sigma_-] + (1 - q)R \left[C_a - B + E - \frac{1}{2} \right],$$

and

$$E(\pi_{d,ns}^q) - E(\pi_d^G) = -\frac{1}{2}VE - \frac{1}{2}q^2V [C_a - B + E - \sigma_-] - qR \left[C_a - B + E - \frac{1}{2} \right].$$

Note that $C_a - B + E - \sigma_- > 0$ by Assumption 2. Thus, regardless of the sign of $[C_a - B + E - \frac{1}{2}]$, either $E(\pi_{nd}^G)$ or $E(\pi_d^G)$ (or both) are greater than $E(\pi_{d,ns}^q)$. Thus, setting $q = 0$ with the appropriate policy rule (discipline or no discipline) dominates setting $q > q^{\max}$ when case (ii) applies.

Case (iii): the government never selects implementation. This strategy obviously yields a lower payoff than $E(\pi_d^G)$. In both cases, in fact, the government is always supported by the public. In the latter case, however, the government selects implementation when the signals are both positive. By Assumption 1, that yields a larger payoff than selecting the status quo.

■

To prove Lemma 2 and Proposition 1, the following lemma is useful.

Lemma B1. *Suppose the government and the public share the same beliefs about the true state: $\sigma_{Gov} = \sigma_P$. Then, whenever the public supports implementation, the government also selects implementation. If the public supports the status quo, the government selects implementation when $\sigma_{Gov} \geq C_a - B + E$.*

Proof of Lemma B1. Let $\sigma_{Gov} = \sigma_P = \sigma$. Recall that the public supports implementation if $\sigma \geq C_a$. Assuming that the public supports implementation, the government selects

implementation if $\sigma \geq C_a - B - E$. This condition is obviously implied by $\sigma \geq C_a$. Thus, whenever the public supports implementation, the government also selects implementation. The second part of the lemma follows immediately from equation (2.2). ■

Proof of Lemma 2. From Assumption 1 and the fact that $q \leq q^{\max}$, it follows immediately that the public supports implementation if and only if both signals are positive. Because the public supports implementation when both signals are positive, it also follows from Lemma B1 that the government must select implementation in that case. Finally, Assumption 2 implies that the government selects the status quo when both signals are negative. ■

Proof of Proposition 1. From Lemma 2, we know that in equilibrium the government selects implementation when both signals are positive and the status quo when both signals are negative. Moreover, the public supports implementation if and only if both signals are positive. The only thing that remains to be shown is whether implementation or the status quo is selected by the government when the evidence is mixed.

Suppose that $B < C_a - \frac{1}{2} + E$ (cases (i) and (ii)). Recall that $\sigma^q = \Pr(A | (\alpha, \emptyset)^q) \leq \frac{1}{2}$. Because $\sigma^q \leq \frac{1}{2} < C_a - B + E$, Lemma B1 implies that for all q 's the government selects the status quo when the evidence is mixed. Thus, if $B < C_a - \frac{1}{2} + E$, for all q 's the government is disciplined by public opinion.

Next, we derive the optimal q when $B < C_a - \frac{1}{2} + E$ (cases (i) and (ii)). Recall that $E(\pi_d^q)$ denotes the government payoff when the government is disciplined by public opinion and the size of the bias is q . Define

$$\Delta\pi_{d,nd}^q \equiv E(\pi_d^q) - E(\pi_{nd}^G), \quad \text{where } q \in [0, q^{\max}].$$

The optimum q solves

$$\max_{q \in [0, q^{\max}]} E(\pi_d^q),$$

or equivalently

$$\max_{q \in [0, q^{\max}]} \Delta\pi_{d,nd}^q,$$

since $E(\pi_{nd}^G)$ is independent of q . Thus the equilibrium level of bias is $q^* = \arg \max_{q \in [0, q^{\max}]} E(\pi_d^q)$.

In Appendix A we showed that

$$\Delta\pi_{d,nd}^q = qRE - \frac{1}{2}q^2V [C_a - B - \sigma_-] + (1 - q)R \left[C_a - B + E - \frac{1}{2} \right],$$

(see equation (A7)). By Assumption 2, $C_a - B - \sigma_- > 0$. Thus simple algebra yields

$$q^* = \begin{cases} \frac{R}{V} \frac{1/2 - C_a + B}{C_a - B - \sigma_-} & \text{if } B \in (C_a - \frac{1}{2}, C_a - \frac{1}{2} + E) \\ 0 & \text{if } B \leq C_a - \frac{1}{2} \end{cases}.$$

Note that given condition (3.2), we restrict attention to parameter values for which the constraint $q \leq q^{\max}$ does not bind. Using the above formula for q^* , we can rewrite condition (3.2) more explicitly as

$$\underbrace{\frac{R}{V} \frac{1/2 - C_a + B}{C_a - B - \sigma_-}}_{q^*} \leq \underbrace{\frac{-\frac{R}{V}(C_a - 1/2) + \sqrt{(\frac{R}{V})^2 (C_a - 1/2)^2 - (C_a - \sigma_-)(C_a - \sigma_+)}}{(C_a - \sigma_+)}}_{q^{\max}}. \quad (\text{B1})$$

Define $\bar{B}_1 \in (C_a - \frac{1}{2}, C_a - \sigma_-)$ as the value of B for which (B1) holds with equality, where (B1) holds strictly for all $B < \bar{B}_1$. Hence there exist values of $B \in (C_a - \frac{1}{2}, C_a - \frac{1}{2} + E)$ that satisfy condition (3.2). (If instead $B \leq C_a - \frac{1}{2}$, then $q^* = 0$ and assumption (3.2) always holds.) This proves parts (i) and (ii) of Proposition 1.

To prove part (iii), let $B \geq C_a - \frac{1}{2} + E$. Two cases can arise, depending on whether q is above or below the threshold \hat{q} . \hat{q} is defined as the minimum between 1 and the solution to $\sigma^q = C_a - B + E$. Recall that $\sigma^q = \frac{\frac{R}{2} + q(1-\theta)^2}{R+qV}$. Then

$$\hat{q} = \frac{R(\frac{1}{2} - C_a + B - E)}{V(C_a - B + E) - (1 - \theta)^2} = \frac{R}{V} \frac{1/2 - C_a + B - E}{C_a - B + E - \sigma_-} \leq 1, \quad (\text{B2})$$

if $C_a - B + E \geq 1 - \theta$, and $\hat{q} = 1$ if $C_a - B + E < 1 - \theta$ (the requirement that $B \geq C_a - \frac{1}{2} + E$ implies $\hat{q} \geq 0$).

Note that when $q < \hat{q}$, then it is optimal for the government to select implementation when the signals are mixed, because $q < \hat{q}$ implies $\sigma^q > C_a - B + E$ (see Lemma B1). Conversely, when $q > \hat{q}$, it is optimal for the government to select the status quo. Thus, when $B \geq C_a - \frac{1}{2} + E$, the normalized government payoff is

$$\Delta\pi_{LB}^q \equiv \begin{cases} \Delta\pi_{nd,nd}^q = E(\pi_{nd}^q) - E(\pi_{nd}^G) & \text{if } q \in [0, \hat{q}] \\ \Delta\pi_{d,nd}^q = E(\pi_d^q) - E(\pi_{nd}^G) & \text{if } q \in (\hat{q}, q^{\max}] \end{cases},$$

(note that $\hat{q} \leq q^{\max}$ is implied by (3.2) since $\hat{q} \leq q^*$). The optimum q solves

$$\max_{q \in [0, q^{\max}]} \Delta\pi_{LB}^q.$$

It is easy to show that $\Delta\pi_{LB}^q$ has the following properties: (i) $\Delta\pi_{nd,nd}^q$ is strictly convex in q , (ii) $\Delta\pi_{d,nd}^q$ is strictly concave in q and achieves its maximum at $q^* = \frac{R}{V} \frac{1/2 - (C_a - B)}{C_a - B - \sigma_-} \geq \hat{q}$, (iii)

$\frac{\partial \Delta \pi_{d,nd}^q}{\partial q} \geq 0$ at $\hat{q} < 1$, and (iv) $\Delta \pi_{LB}^q$ is continuous on $[0, q^{\max}]$. Indeed, simple computations yield $\frac{\partial^2 \Delta \pi_{nd,nd}^q}{\partial q^2} = V(C_a - B + E - \sigma_- + E) > 0$ and $\frac{\partial^2 \Delta \pi_{d,nd}^q}{\partial q^2} = -V(C_a - B - \sigma_-) < 0$.

Moreover

$$q^* \equiv \arg \max_{q \in [0,1]} \Delta \pi_{d,nd}^q = \frac{R/2 - C_a + B}{V(C_a - B - \sigma_-)}. \quad (\text{B3})$$

Note that since $q^* \geq \hat{q}$ and $\Delta \pi_{d,nd}^q$ is strictly concave, $\left. \frac{\partial \Delta \pi_{d,nd}^q}{\partial q} \right|_{q=\hat{q}} \geq 0$. Finally, to show that $\Delta \pi_{LB}^q$ is continuous on $[0, q^{\max}]$, note that $\Delta \pi_{nd,nd}^q$ and $\Delta \pi_{d,nd}^q$ are both continuous in q on their respective domains. Using (B2), it is also easy to verify that $\Delta \pi_{nd,nd}^q = \Delta \pi_{d,nd}^q$ at \hat{q} .

Proposition 1(iii) follows from these properties of $\Delta \pi_{LB}^q$. Note in fact that because $\Delta \pi_{nd,nd}^q$ is strictly convex, then $\arg \max_{q \in [0, \hat{q}]} \Delta \pi_{nd,nd}^q$ is either 0 or \hat{q} . However, \hat{q} yields a lower payoff than q^* since $\Delta \pi_{nd,nd}^{\hat{q}} = \Delta \pi_{d,nd}^{\hat{q}} \leq \Delta \pi_{d,nd}^{q^*}$.²³ Thus, on $[0, q^{\max}]$, the optimal q is either 0 or q^* , depending on whether $E(\pi_{nd}^0) \leq E(\pi_d^{q^*})$ (or, equivalently, $\Delta \pi_{nd,nd}^0 \leq \Delta \pi_{d,nd}^{q^*}$). Finally, since $\bar{B}_1 > C_a - 1/2$, there exist values of $B > C_a - 1/2 + E$ that satisfy condition (3.2), provided that E is sufficiently small. ■

Proof of Proposition 2. Suppose first that $B \leq C_a - \frac{1}{2}$. Then by Proposition 1(i), the equilibrium level of bias is zero. Imposing $q = 0$ therefore leaves the public's payoff unchanged.

Suppose instead that $B \in (C_a - \frac{1}{2}, C_a - \frac{1}{2} + E)$. Then by Proposition 1(ii), the government is disciplined by public opinion and the equilibrium level of bias is $q^* = \frac{R/2 - C_a + B}{V(C_a - B - \sigma_-)}$. If we impose $q = 0$, (A7) implies the government will remain disciplined by public opinion since

$$\Delta \pi_{d,nd}^G = E(\pi_d^G) - E(\pi_{nd}^G) = R \left(C_a - B + E - \frac{1}{2} \right) > 0.$$

Let $E(U_d^q)$ denote the public's payoff under discipline and with bias q (see equation (A10)).

We have

$$\Delta U_{d,d}^q \equiv E(U_d^q) - E(U_d^G) = -qR \left(C_a - \frac{1}{2} \right) - \frac{1}{2} q^2 V(C_a - \sigma_-),$$

which is decreasing in q . Hence imposing $q = 0$ when $B \in (C_a - \frac{1}{2}, C_a - \frac{1}{2} + E)$ will strictly increase the public's payoff.

Now suppose that $B \geq C_a - \frac{1}{2} + E$. Proposition 1(iii) then implies that for $\Delta \pi_{d,nd}^{q^*} \leq 0$, the equilibrium level of bias is zero, so that imposing $q = 0$ has no effect. If instead $\Delta \pi_{d,nd}^{q^*} > 0$, then the government is disciplined by public opinion and the equilibrium level of bias is q^* .

²³ Actually, because $B > 0$ in case (iii), $\Delta \pi_{d,nd}^{\hat{q}} < \Delta \pi_{d,nd}^{q^*}$.

Moreover, by (A7), $B \geq C_a - \frac{1}{2} + E$ implies $\Delta\pi_{d,nd}^G \leq 0$. Hence imposing $q = 0$ leaves the government undisciplined by public opinion. By (A13), this will strictly decrease the public's payoff if

$$\Delta U_{d,nd}^{q^*} = (1 - q^*) R \left(C_a - \frac{1}{2} \right) - \frac{1}{2} q^{*2} V (C_a - \sigma_-) > 0.$$

Direct substitution yields $\Delta U_{d,nd}^G > 0$ and $\Delta U_{d,nd}^1 < 0$, where $\Delta U_{d,nd}^{q^*}$ is decreasing in q^* . Moreover, $q^* = \frac{R}{V} \frac{1/2 - C_a + B}{C_a - B - \sigma_-}$ is increasing in B , with $q^* = 0$ when evaluated at $B = C_a - \frac{1}{2}$. Define $\bar{B}_2 > C_a - \frac{1}{2}$ as the value of B for which $\Delta U_{d,nd}^{q^*} = 0$, where $\Delta U_{d,nd}^{q^*} > 0$ if and only if $B < \bar{B}_2$. Moreover, define $\bar{B} = \min(\bar{B}_1, \bar{B}_2)$, where (B1) holds strictly if and only if $B < \bar{B}_1$. It follows that, over the parameter region for which condition (3.2) does not bind, imposing $q = 0$ will strictly decrease the public's payoff if and only if $B \in [C_a - \frac{1}{2} + E, \bar{B})$ and

$$\Delta\pi_{d,nd}^{q^*} = q^* R E - \frac{1}{2} q^{*2} V [C_a - B - \sigma_-] + (1 - q^*) R \left[C_a - B + E - \frac{1}{2} \right] > 0.$$

Fix $B > C_a - \frac{1}{2}$. Note that $q^* = \frac{R}{V} \frac{1/2 - C_a + B}{C_a - B - \sigma_-}$ is independent of E , $\Delta\pi_{d,nd}^{q^*}$ is strictly increasing in E , and $\Delta\pi_{d,nd}^{q^*} < 0$ when evaluated at $E = 0$. Define $\underline{E}(B)$ as the value of E for which $\Delta\pi_{d,nd}^{q^*} = 0$:

$$\underline{E}(B) = \frac{\frac{1}{2} q^{*2} V [C_a - B - \sigma_-] + (1 - q^*) R [B - C_a + \frac{1}{2}]}{R},$$

so that $\Delta\pi_{d,nd}^{q^*} \leq 0$ for all $E \leq \underline{E}(B)$, and $\Delta\pi_{d,nd}^{q^*} > 0$ for all $(\underline{E}(B), B - C_a + \frac{1}{2}]$. To complete the proof, it remains to show that $\underline{E}(B) < B - C_a + \frac{1}{2}$. This is the case since $\underline{E}(B)$ is decreasing in q^* over the interval $[0, \frac{R}{V} \frac{1/2 - C_a + B}{C_a - B - \sigma_-}]$, and $\underline{E}(B) = B - C_a + \frac{1}{2}$ when evaluated at $q^* = 0$. ■

Proof of Proposition 3. Standard results in decision theory imply that, for any *given* belief about q that the public may hold, setting $q = 0$ is a weakly dominant strategy for the government (see Marschak and Radner, 1972, pp. 65-67. Further details are available from the authors upon request).

To prove Proposition 3, therefore, we set $q = 0$ and check that the posited strategies form an equilibrium. For brevity's sake, we focus on case (ii), which is the most interesting case since it involves mixed strategies. Checking that the posited strategies form an equilibrium for all the other cases is straightforward.

Let $B \in (C_a - \frac{1}{2}, C_a - \frac{1}{2} + E)$ and $\hat{\sigma} < C_a$. Then the government always selects the status quo when $\mathbf{s}^G = (\emptyset, \emptyset)$. Indeed,

$$-C_n \sigma_- > -C_a (1 - \sigma_-) + B - \check{z} E \iff C_a - B > \sigma_- - \check{z} E.$$

It is also straightforward to check that the public always supports the status quo when $p = n$.

Note that, if the government always selects implementation when the evidence is mixed, then the public must support the status quo when $p = a$ since $\hat{\sigma} < C_a$. (Recall that $\hat{\sigma}$ is defined as the public's belief that A is the true state when $q = 0$, $p = a$, and a is selected by the government if and only if $\mathbf{s}^G \neq (\emptyset, \emptyset)$.) But this implies that it cannot be optimal for the government to always implement the project when the evidence is mixed because, by deviating and selecting the status quo, it would get an higher payoff:

$$-C_a\left(1 - \frac{1}{2}\right) + B - E < -C_n\frac{1}{2} \iff B < C_a - \frac{1}{2} + E.$$

Similarly, if the government always selects the status quo when the evidence is mixed, then the public must always support implementation when $p = a$ since $\sigma_+^G > C_a$ (by Assumption 1). But then it is not optimal for the government to always select the status quo since $B > C_a - \frac{1}{2}$. Thus the equilibrium must be in mixed strategies.

To characterize the mixed strategy equilibrium, we need to find a probability \check{z} that makes the government indifferent between selecting a and n when the evidence is mixed, and a probability \check{s} that makes the public indifferent between supporting a and n when $p = a$. The indifference condition for the government is given by

$$-C_n\frac{1}{2} = -C_a\left(1 - \frac{1}{2}\right) + B - \check{z}E \implies \check{z} = \frac{1/2 - (C_a - B)}{E}.$$

The indifference condition for the public is $\Pr(A | a, \check{s}) = C_a$, where $\Pr(A | a, \check{s})$ is the public's belief that $S = A$ when the government selects implementation ($p = a$), given that the government implements the project with probability one if $(\alpha, \alpha)^G$, with probability \check{s} if the evidence is mixed, and with probability zero if $(\emptyset, \emptyset)^G$. Using Bayes' rule

$$\Pr(A | a, \check{s}) = \frac{\theta^2 + \check{s}R}{V + 2\check{s}R} = C_a \implies \check{s} = \frac{\theta^2 - C_a V}{R(2C_a - 1)} = \frac{\sigma_+^G - C_a}{2r(C_a - 1/2)}.$$

It is easy to show that $\hat{\sigma} < C_a$ implies $\check{s} < 1$ and that $B \in (C_a - \frac{1}{2}, C_a - \frac{1}{2} + E)$ implies $\check{z} < 1$.

Finally, note that if the government selects implementation with positive probability when the signals are mixed, then it must select implementation with probability one when the signals are both positive. ■

Proof of Proposition 4. Obvious. ■

Proof of Proposition 5. Note that Constitutions III and IV yield the same citizen welfare since nondisclosure implies $q = 0$ in equilibrium (Proposition 3). Moreover, by Proposition 4

Constitution I dominates Constitution III (and hence also Constitution IV). Thus disclosure is always a feature of an optimal constitution. Finally, Propositions 1 and 2 demonstrate that either Constitution I or II can be optimal. ■