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

Andrea Patacconi

University of Oxford

February 2007

Abstract: Policymakers often motivate their decisions by disclosing information. While this can help hold the government to account, it may also give incumbents an incentive to "fix the evidence" around their preferred policy. This paper studies how important elements of institutional design affect citizen welfare when such incentives are present. When the process of information gathering and evaluation can be easily distorted by an incumbent executive ('nonindependent agencies'), making government more transparent creates a tradeoff between manipulation of information and accountability. Disclosure of information always benefits the public when government agencies can be insulated from political pressure. However, nonindependent agencies may be part of an optimal institutional arrangement.

KEYWORDS: Open Government, Accountability, Independent Agencies, Manipulation of Information, Intelligence Failures.

JEL CLASSIFICATION: D73, H11, H56.

---

[1]Postal address: Andrea Patacconi, Department of Economics, Manor Road Building, Manor Road, Oxford, OX1 3UQ, United Kingdom. E-mail: andrea.patacconi@economics.ox.ac.uk.

I would like to thank David Austen-Smith, Sharon Belenzon, Karen Croxson, Diego Gambetta, Ian Jewitt, Kohei Kawamura, David Myatt, Torsten Persson, John Quah, Dimitri Vayanos, Jeffrey Zwiebel, seminar participants at Bristol, Oxford and Salerno, and especially Clare Leaver, Meg Meyer and Andrea Prat for helpful discussions and remarks. Financial support from the Economics Department at the University of Oxford, the Royal Economic Society and the British Academy is gratefully acknowledged. All remaining errors are mine.

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

Government accountability is an essential element of democracy. Policymakers in democratic countries must frequently explain their actions and policies and can be sanctioned if they misuse their power. In the US, for instance, the Constitution (Article I, Section 5) requires that each house of Congress "keep a Journal of its Proceedings, and from time to time to publish the same, excepting such Parts that may in their judgement require Secrecy", and similarly the president, whenever he vetoes a bill, is required to state his objections which must then be published in the journal of the house in which the bill originated. Indeed, it is clear that only with timely and accurate information about what the office-holders are doing, can accountability mechanisms such as the electoral process work properly.

This paper explores a dark side of transparency in government. The key idea is that, if information in the possession of public bodies is disclosed, policymakers may have the incentive to distort the process of information gathering and evaluation so that information could be used to sell government policies. Such incentives naturally arise when policymakers are concerned about their reelection prospects and may compromise the quality of the decision-making process, as the following examples suggest.

Open Government Laws. The principle of open government is spreading around the world. Banisar (2004) reports that in 2004 over fifty countries, including the US and most EU members, had adopted Freedom of Information (FOI) laws, with over thirty more having pending efforts. Coexisting with a positive view of transparency, however, there are also concerns that opening government to public scrutiny might compromise the way in which information is gathered and disseminated. According to observers of EU politics, for instance, "The Economic and Financial Committee is characterized by frank exchanges of views about economic policy, precisely because it is a more insulated setting than is ECOFIN, the Council of Ministers of Finance" (Stasavage (2006) p.15). Similarly, FOI experts have highlighted several bureaucratic practices, including changes in recordkeeping, decline in candor, manipulation and failure to create records, that have been used to undercut the right to information.[1] As those practices might impair decision-making, most FOI laws provide some form of protection for the exchange of views that takes place within government. Typically, such protection is granted by exempting pre-decision information from the law.[2]

Targets and Government Figures. A distinctive feature of British politics under New Labour is the proliferation of performance targets for the public sector. The aim of such targets is to provide a benchmark against which public sector output can be measured. The problem is that once the credibility of a government is tied to the official figures that measure its performance against the targets, the pressure to produce favorable statistics may become irresistible.[3] This may help explain why only 26% of accident and emergency (A&E) departments thought the figures they submitted for the waiting-time target were a true reflection of their performance, and 16% reported direct manipulation of data (British Medical Association, 2005). Many Britons also believe that official statistics are 'politicized'. A recent survey by the Office for National Statistics in Britain found that the majority of respondents believed official figures were changed to support a particular argument (68%), that there was political interference in their production (58%), and that mistakes were suppressed (69%) (Jones and Jones (2005)). The Statistics Commission (2005) highlights two reasons why a lack of trust in official statistics should concern the public. First, "official statistics influence large numbers of 'business' decisions day-to-day, across government, the public services and in the private and not-for-profit sectors" (p.7). Secondly, trust in official figures must be

[1] See Roberts (2006). As a concrete example, he reports that, to limit the damage done by the release of information under the law and acting on instructions from the Prime Minister's office, the Canadian Department of Public Works kept "minimum information" on the spending of a sponsorship program and later created expenditure guidelines for cosmetic purposes.

[2] However, the degree to which pre-decision material is protected varies considerably across countries. At one extreme, in fact, "The Danish law has a wide class exemption for internal case material, including documents prepared by the authority for its own use. These remain exempt after decisions are taken" (Frankel 2001, p.8). In Sweden, instead, pre-decision information must be disclosed after a decision has been "finally settled".

[3] As The Economist puts it, in fact, "The more that ministerial reputations ride on statistics, the more protection statisticians need from political interference. Current arrangements fail to deliver that protection." (26/03/05, p.31-32).

preserved because "if the general public do not trust the statistics that underpin and explain policy initiatives, and that measure the success of those policies, then government and public bodies themselves will not be trusted" (p.4). Both these issues will figure prominently in this paper.

Intelligence Failures. The last few years have witnessed an explosion of interest in the workings of the intelligence community. One reasons has been the tremendous hit taken by America and Britain's intelligence services when Iraq's expected stock of weapons of mass destruction (WMD) failed to materialize. A major controversy revolves around whether the intelligence community judgements were warped by inappropriate political pressure. Policymakers were obviously aware of the importance of intelligence in shaping public opinion.[4] There is also, as The Economist puts it, "a good deal to suggest that the [American] administration employed strategies to mould the intelligence to its purposes" (17/07/04, p.25).[5] In particular, it appears that some senior intelligence analysts engaged in unprofessional behaviors to please their superiors. Take the case of Curve Ball, a drunkard whose claims constituted the backbone of the intelligence on Iraq's mobile biological weapons program mentioned by the U.S. Secretary of State Colin Powell in February 2003 at the U.N. Security Council. When, before Mr. Powell delivered his speech, 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".[6] As a result of the Deputy Chief's (in)action, potentially useful information was not transmitted to top policymakers, thus perhaps contributing to biased decision-making.[7]

[4] Referring to JIC dossier of September 2002, for instance, Lord Butler writes that "The Government wanted an unclassified document on which it could draw in its *advocacy* of its policy" (p.154, emphasis in original).

[5] See The Economist's 'Special Report Intelligence Failures' (17/07/04, pp.23-25) for a comprehensive discussion. It is also important to note that the UN inspectors immediately raised serious concerns about the quality of the American and British intelligence (see Blix (2005), especially pp.156-157 and 175-178). On the other hand, according to the WMD Commission, no analytic judgement had been made or changed by the intelligence community in response to political pressure to reach a particular conclusion.

[6] Extract from an e-mail provided to the WMD Commission, p.249.

[7] See Prendergast (1993) and Garicano and Posner (2006). Career concerns may also exacerbate 'yes-man' behavior. In that respect, it is not encouraging 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" (WMD Commission, p.193).

This paper develops a simple political agency model to highlight a key tradeoff between transparency and manipulation of information. As in every agency model, the preferences of the agent (the government) and those of the principal (the public) are not perfectly aligned. We assume that (i) the government wants to be perceived as taking the 'right' decision by the public, and (ii) the process of information gathering and evaluation can be manipulated by the government. Thus, we implicitly posit that the government agencies in charge of collecting information are not independent in that they are not fully insulated from political pressure (but see below). In a setting where pre-decision information is disclosed we show how the 'politicization' of information can be used to boost support for the government policy despite the fact that citizens are sophisticated and thus skeptical about the information they receive. We then use the model to investigate how different institutional environments ('constitutions') affect citizen welfare. Of particular interest is the question of whether secrecy – a commitment not to disclose information – can protect the integrity of the decision-making process. In defending the principle behind executive privilege, for instance, the US Supreme Court stated that the "need for protection of communications between high Government and those who advise and assist them in the performance of their manifold duties [...] is too plain to require further discussion" (US v. Nixon (1974)). Furthermore, as mentioned above, one of the most common exemptions to the principle of open government in practice is the exemption of pre-decision information. We highlight some of the costs (in terms of manipulation of information) and benefits (in terms of greater accountability) of transparency in government and evaluate how different disclosure environments affect citizen welfare. We also highlight some non-obvious implications of fully insulating government agencies from political pressure.

1.1 Overview

We now provide an overview of the paper and the main results. In Section 2, we set out the basic model. There are two players, a government (policymakers, etc.) and the public (citizens, the median voter, etc.). The government must choose whether or not to implement a policy. The outcome of that choice depends on an unknown state of the world, but the government also enjoys a private benefit when the policy is implemented. This creates a potential conflict of interest between the government and the public. Before taking its decision, the government collects information about the state of the world.

Citizens play no explicit role in the decision-making process. Nevertheless, their opinion matters because the government wants to be perceived as taking the 'right' decision by the public, say because it wants to improve its political standing and/or electoral chances. In this way we capture the disciplining effect of public opinion on policy and measure the government's electoral (or legitimacy) concerns. The public decides whether or not to support the government policy on the basis of the available information, which may include information disclosed by the government. A key assumption of the model is that policymakers can distort the process of information gathering and evaluation, should that be in their interest. More precisely, we posit that the government may allow fake signals that support its initial disposition (relative to the public) percolate through the screening process. Examples include the government seeking the advice of experts who are known to be biased in favor of a particular policy or (explicitly or implicitly) encouraging biased information gathering and evaluation. In line with second interpretation, we will sometimes say that the government can bias the 'screening technology' and will refer to the extent to which information is manipulated as the 'size of the bias'. The drawback we emphasize is that, in an attempt to justify a certain course of action, the quality of the information brought to bear on a decision may be compromised.

Information can be disclosed to the public but once potentially biased information is gathered, it must be reported truthfully, if reported at all.[8] We focus on two polar scenarios, one in which the government discloses all the information and another in which the government commits not to disclose any information. In the full disclosure scenario, we further distinguish between situations where the size of the bias in information is observed by the public (Section 3) and the case when the size of the bias is not observable and is only inferred in equilibrium (Section 4).

We show that the equilibria of the full disclosure model with observable bias and with unobservable bias are qualitatively similar. In both cases, policymakers face a tradeoff between preserving the quality of the information brought to bear on the decision and the benefits of distorting the information made available to the public. Biasing the information can be optimal because the government wants to maximizes the likelihood that its chosen

[8] One might think of situations where policymakers can be severely punished if caught lying, but blame can be placed on bureaucrats (or informants) should the information turn out to be forged.

policy is supported by the public. By biasing the screening technology appropriately, the government can increase the likelihood that signals supporting its initial disposition will be put into the public domain. Public confidence on the merits of the various proposals will be low, but if the bias in information is not too large citizens will be more likely to support the government policy. However, the government must preserve a minimum level of trust in official communications, without which the public would pay no attention to them.

Section 5 studies a nondisclosure scenario where the government commits not to disclose any information, the size of the bias in information is unobservable and the public updates beliefs by observing the government's choice. In line with conventional wisdom, we show that the government has no incentive to manipulate the information. However, since the public cannot figure out exactly why a specific decision was taken, the government is less accountable and thus less responsive to public desires, relative to the full disclosure scenario. Thus either scenario can be optimal from the public's viewpoint.

Section 6 examines the implications of granting independence to the government agencies in charge of collecting and evaluating pre-decision information (modelled here as a commitment *not* to manipulate the information). We show that if government agencies are independent, disclosure of information always dominates nondisclosure from the citizens' viewpoint. We then study optimal constitutional design along both dimensions: transparency vs. secrecy and independent vs. nonindependent agencies. Transparency is always a feature of an optimal constitution. However, an optimal design may require nonindependent agencies. Surprisingly, in fact, disclosure sometimes benefits the public exactly because it creates incentives for manipulation. This is because manipulation of information creates two opposite effects. On the one hand, biased information hurts the public because the wrong decision may be taken. On the other hand, concerns about forgeries induce policymakers to implement the project less often, which is what the public wants. We show that this second, beneficial effect can dominate the first.

Section 7 discusses a number of extensions. First, we argue that as in most models with hard information, voluntary disclosure will typically result in all the information being disclosed. We then explore the idea that public opinion might be more easily manipulated in foreign policy than in domestic policy (Shapiro and Page (1992)). We show that small differences in the likelihood that forged signals are detected can drastically change the equilibrium level of manipulation. Finally, we investigate whether it can ever be beneficial for the government to bias the process of information gathering and evaluation *against* its initial disposition. Our answer will critically hinge upon whether the size of the bias in information is observable or not. Section 8 concludes. All the proofs are gathered in two technical Appendixes.

1.2 Related Literature

This paper directly contributes to a growing literature focusing on transparency in agency relationships. In a classic paper, Holmstrom (1979) shows that more information is typically beneficial to a principal as it allows to better assess an agent's performance and discipline him. Exceptions to the Holmstrom's result are provided by Cremer (1995) and Holmstrom (1999), among others.[9] Stasavage (2004) also points out that in negotiations the possibility of closed-door meetings may facilitate attempts to strike bargains. However, the paper most closely related to the present one is Prat (2005). Prat explores the idea that a certain kind of transparency may create perverse incentives. He develops a model of career concerns for experts where a principal may observe an agent's action and/or the consequences of that action. His main result is that transparency on action may induce the agent to disregard useful private information and act in a conformist manner. As a result, the principal may be better off by committing not to observe the action. By contrast, transparency on consequences is always beneficial to the principal.

The present paper focuses neither on transparency on action nor on consequences. In this model transparency measures the extent to which pre-decision information is shared between the agent and the principal. Actions are observable but their long-term consequences are not (these are reasonable assumptions in politics and could be relaxed). The key issue is not whether transparency induces conformism on the part of the agent, but whether an agent would distort his own information (and possibly the principal's) to influence how the principal perceives his action.[10] In contrast to Prat, in fact, there is no 'smart' action in this model (a smart action being the one that an able agent is expected to choose a priori). Here

[9] Prat (2006) lists several other reasons for why transparency may not always be desirable, including the right to privacy, the direct cost of disclosure and the risk that hostile parties learn sensitive information.

[10] Maggi and Rodriguez-Clare (1995) and Dewatripont and Tirole (1999) also look at the issue of manipulation of information in agency but the focus of their analysis is very different.

the government looks good when its chosen policy is believed to match the true state of the world. This can be accomplished in two ways. One is to select the policy that in the light of the best available evidence is more likely to be optimal from the public's viewpoint. The dysfunctional alternative we highlight is to choose an a priori favored policy and then to "fix the evidence" around it.[11]

This paper also contributes to a burgeoning literature on accountability in government (see, e.g., Persson, Roland and Tabellini (1997), Canes-Wrone, Herron and Shotts (2001), Besley and Burgess (2001, 2002) and Maskin and Tirole (2004)).[12] Particularly relevant is a recent paper by Besley and Prat (2006) who focus on the relationship between the government and the media. In their model incumbents can manipulate the transmission of information to the public by offering some form of compensation to the media owners. Their focus is on how features of the media industry such as concentration and ownership affect political outcomes. Here instead the focus is on how the production of information that can be distorted by incumbents. As a result, not only can public opinion be manipulated, but policy-making may also be distorted. Our main contribution is to show how the transparency and independence of government agencies affect both the quality of the information brought to bear on a decision and government accountability.

2 The Model

Preferences. There are two players, the government and the public. The government gathers some information and then decides whether or not to implement a project. Let $p \in \{w, n\}$ be the government policy, where w stands for implementation and n denotes the status quo. Citizens observe p and any disclosed information and then choose which policy $v \in \{w, n\}$ they support. Their choice is made before outcomes are observed.

Players maximize expected payoffs. These payoffs depend on the true state of the world, $S \in \{W, N\} \equiv \Lambda$, which is initially unknown. Prior beliefs about the state of the world sum-

[11] In this respect, the present paper also differs from recent contributions focusing on transparency in group decision making (see, e.g., Levy (2005)). The intuition there is similar to Prat's: Secrecy may induce better decisions because, if individual votes cannot be observed, the members of a committee have less of an incentive to distort their actions in order to signal their types. See also Swank, Swank and Visser (2006).

[12] Relatedly, Biglaiser and Mezzetti (1997), Milbourn, Shockley and Thakor (2001) and Suurmond, Swank and Visser (2004) study how reputational concens affect the incentives to implement projects and invest in information.

marize all the information that is publicly available before the government gathers additional information. However, for simplicity we assume the states to be a priori equally likely.[13] The public wants the policy to match the true state of the world. For instance, citizens may want the government to build new nuclear plants if they believe these plants to be reasonably safe, but not otherwise. Formally, the public is assumed to incur a loss of C_w when the chosen policy is w but the true state is N, and a loss of C_n when $p = n$ but $S = W$. The payoff when the policy matches the true state of the world is normalized to zero. With no loss of generality C_w and C_n are also normalized so that $C_w + C_n = 1$.

The public's payoff is determined solely by the above considerations and therefore does not depend on v, the policy they support. Citizens choose v as follows. Let σ_P denote the public's posterior belief that $S = W$. (The information structure of the game is discussed later.) Citizens support the action that maximizes their expected payoff. In particular, they support w if

$$-C_w (1 - \sigma_P) \geq -C_n \sigma_P \quad \Leftrightarrow \quad \sigma_P \geq C_w \equiv T_P \tag{1}$$

(for now, we arbitrarily break ties in favor of w). Similarly, the public supports n if $\sigma_P < T_P$.

The preferences of the government are more complex. Realistically, the government cares about citizen welfare (a 'legacy' concern). However, the government also enjoys a private benefit $B \geq 0$ when the project is implemented (e.g., bribes, an ideological bias, etc.). This creates a potential conflict of interest between the government and the public. Finally, to capture the disciplining effect of public opinion, the government is assumed to incur a loss $E \geq 0$ whenever its chosen policy is not supported by the public. Different interpretations can be attached to this cost. The most straightforward one is in terms of reelection prospects: If a government adopts an unpopular policy, citizens may vote for the opposition in the next election.[14] Alternatively, E could measure several other costs linked to a loss of popularity: public vilification by the press, reduced job opportunities in the private sector, strong manifestations of protest, etc. The notion of "public" should also be interpreted broadly. The model could apply for instance to scenarios where a country wants to convince other nations that a particular course of action is justified to receive logistic or military support. Here E would measure the loss to that country should support be denied.

[13] This assumption could be easily relaxed.

[14] Note however that since citizens are assumed to make up their minds and vote before outcomes are observed, the government policy should be interpreted as a 'long-term' project.

These ideas are formalized as follows. Let σ_{Gov} denote the government's belief that $S = W$. Let $1_{\{d,v\}}$ be an indicator function taking value 1 if the government policy is not supported by the public ($p \neq v$), and 0 otherwise. The government implements the project if

$$-C_w[1-\sigma_{Gov}] + B - 1_{\{w,v\}}E \geq -C_n\sigma_{Gov} - 1_{\{n,v\}}E$$

$$\Leftrightarrow \qquad \sigma_{Gov} \geq C_w - B + (1_{\{w,v\}} - 1_{\{n,v\}})E. \tag{2}$$

Thus the project is implemented when the probability that W is the true state is sufficiently high. Relative to (1), however, the government's threshold tends to be lower (and the decision more biased toward implementation) because of B. Moreover, the government will also tend to choose whichever policy happens to be 'popular'.

A particularly important case is when the government and the public share the same beliefs.

Lemma 1 *Suppose $\sigma_{Gov} = \sigma_P$. The government always chooses w when the public supports w. However, when the public supports n, the government will choose w if $\sigma_{Gov} \geq T_{Gov} \equiv C_w - B + E$.*

The intuition for this result is straightforward. Since the government is more disposed toward the project than the public is and they share the same information, if there is enough evidence to convince the public that w is the 'right' action, there is also enough evidence to convince the government. The opposite is of course not necessarily true, because of the private benefit B. Obviously, the government will be less likely to select an unpopular policy when E is large. In particular

Lemma 2 *If the parties share the same information and $E \geq B$, the government will always choose the policy that the public supports.*

Information Structure. Before taking a decision, the government receives two signals, $s_i \in \{\alpha, \varnothing\} \equiv \Omega$, $i = 1, 2$. A 'positive' signal α provides evidence in support of W and suggests that the project should be implemented. A 'negative' $\varnothing$-signal suggests that the status quo should be preferred. The signal-generating process can be distorted by the government. The specific idea we explore is that the government may let fake α-signals percolate through the system, thus capturing in a simple way the idea of asymmetric vetting.

More formally, consider two 'genuine' signals, denoted by s_i^G. Genuine signals are informative and conditionally independent. In particular, $\Pr(s_i^G = \alpha|W) = \Pr(s_i^G = \varnothing|N) = \theta$, where $\theta \in (\frac{1}{2}, 1)$ measures the precision of the signal. Before these signals are observed, genuine $\varnothing$-signals are transformed into fake α-signals. The probability that a genuine $\varnothing$-signal is transformed into a fake α-signal is $q \in [0, 1]$, and forgeries are independent. Formally, $\Pr(s_i^q = \alpha \mid s_i^G = \varnothing) = q$, where $\mathbf{s}^q = \{s_1^q, s_2^q\}$ denote the observed signals (clearly $s_i^0 = s_i^G$). The probability that a genuine α-signal is transformed into a fake $\varnothing$-signal is zero (but see Subsection 7.3). Thus, q measures how biased the process of information gathering and evaluation is in favor of the government's initial disposition, relative to the public. The 'size of the bias' q is chosen by the government to maximize its expected payoff, with all q's being equally costly. The latter assumption allows us to focus on the incentives to manipulate the information for purely electoral reasons.[15]

This information structure is can be interpreted in several ways. Consider for instance a government seeking the advice of an expert known to be biased in favor of the project. Suppose the expert succeeds in forging an α-signal when he actually got a $\varnothing$-signal with probability q. Here q measures here how biased the expert is. Alternatively, consider a scenario where signals might have been forged and an agency must screen them. Here q would measure the extent to which signals are asymmetrically vetted, that is, the extent to which the agency is willing to question information that supports the government's initial disposition.

Before proceeding, we should emphasize that the government only observes $\mathbf{s}^q$, not $\mathbf{s}^G$, unless it picks $q = 0$. Thus, if the government observes $s^q = \alpha$, $q > 0$ and reports α, it might be reporting forged information, but it would not be telling a lie. In this sense we say that disclosure is 'truthful'. It is also worth mentioning that our notion of manipulation is an instance of 'garbling' of information since

$$\Pr(\mathbf{s}^q, \mathbf{s}^G \mid S) = \Pr(\mathbf{s}^q \mid \mathbf{s}^G) \Pr(\mathbf{s}^G \mid S) \tag{3}$$

for all $\mathbf{s}^G, \mathbf{s}^q \in \Omega^2$ and $S \in \Lambda$. In other words, the conditional probability distribution $\Pr(\mathbf{s}^q \mid \mathbf{s}^G, S)$ is independent of S, capturing the idea that $\mathbf{s}^q$ is determined solely by $\mathbf{s}^G$, but

[15] Note also that in this model the public does not penalize the government just for distorting the screening technology. This is a strong assumption, but could be easily relaxed.

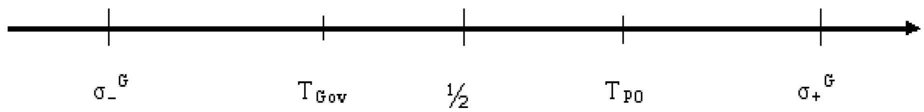

Figure 1: Weak Electoral Concerns: $T_{Gov} \leq \frac{1}{2}$

possibly with the intervention of "noise". To simplify notation, we write $(\cdot,\cdot)^q$ for $\mathbf{s}^q = (\cdot,\cdot)$. Note also that $\Pr(W|(\alpha,\varnothing)^q) = \Pr(W|(\varnothing,\alpha)^q)$ for all q. We define

$$\sigma_+^G \equiv \Pr(W|(\alpha,\alpha)^G) = \frac{\theta^2}{V}, \quad \sigma^G \equiv \Pr(W|(\alpha,\varnothing)^G) = \frac{1}{2}, \quad \sigma_-^G \equiv \Pr(W|(\varnothing,\varnothing)^G) = \frac{(1-\theta)^2}{V}$$

where $V \equiv \theta^2 + (1-\theta)^2 = \Pr((\alpha,\alpha)^G) + \Pr((\varnothing,\varnothing)^G)$ and (for future reference) $R \equiv 2\theta(1-\theta) = \Pr((\alpha,\varnothing)^G) + \Pr((\varnothing,\alpha)^G)$. The following assumptions will be maintained throughout the paper.

Assumption 1 $T_P \equiv C_w \in (\frac{1}{2}, \sigma_+^G]$

Assumption 2 $C_w - B > \sigma_-$

Assumption 1 states that the public supports w when the evidence in favor of W is strong ($\mathbf{s}^G = (\alpha,\alpha)$) but not when it is "mixed" (e.g., $(\alpha,\varnothing)^G$). Assumption 2 can be interpreted as a weak form of congruency between the preferences of the government and those of the public. It says that even without the disciplining effect of public opinion, the government would choose the status quo if the evidence pointed clearly in that direction: $\mathbf{s}^G = (\varnothing,\varnothing)^G$. Thus, in the absence of manipulation, the preferences of the government and the public can only differ when the signals are mixed. This feature of the model motivates the following definition.

Definition. *The government is said to be disciplined by public opinion if, for given q, it implements the project when the signals are $(\alpha,\alpha)^q$ and selects the status quo otherwise. Conversely, the government is said not to be disciplined by public opinion if it implements the project when $s^q \neq (\varnothing,\varnothing)$ and selects the status quo otherwise.*

3 Full Disclosure with Observable Bias

In a full disclosure scenario the government must truthfully disclose its signals. We begin with the case where the size of the bias q is observed by the public. This could be the case because not only the reports but also the identity of the experts (and hence their reputations) are disclosed. The case when q is not observed by the public is studied later.

If the government and the public observe the same signals and q is observable, their common posterior beliefs will be given by

$$\sigma_+^q \equiv \Pr(W|(\alpha,\alpha)^q), \qquad \sigma^q \equiv \Pr(W|(\alpha,\varnothing)^q) = \Pr(W|(\varnothing,\alpha)^k), \qquad \sigma_-^q \equiv \Pr(W|(\varnothing,\varnothing)^q).$$

Note that rational agents tend to discount α-signals more heavily than $\varnothing$-signals since the former might have been forged: $\sigma^q \leq \frac{1}{2}$.[16] Furthermore, the α-signals become less and less informative as the screening technology deteriorates: $\frac{\partial \sigma_+^q}{\partial q}, \frac{\partial \sigma^q}{\partial q} < 0$. Since $\sigma_-^q = \sigma_-^G$ for all q, we will simplify notation by writing $\sigma_-^q = \sigma_-$ henceforth.

Let $\pi(p, v \mid S)$ denote the government's payoff when the policy is p, the public supports v and the state is S. For fixed $q \in [0,1]$, the government's expected payoff is

$$E(\pi^q) = \sum_{\mathbf{s}^q \in \Omega^2} \Pr(\mathbf{s}^q) \underbrace{\sum_{S \in \Lambda} \pi(p(\mathbf{s}^q), v(\mathbf{s}^q) \mid S) \Pr(S \mid \mathbf{s}^q)}_{\pi(p(\mathbf{s}^q), v(\mathbf{s}^q)|\mathbf{s}^q)} \tag{4}$$

where $p(\mathbf{s}^q)$ and $v(\mathbf{s}^q)$ denote, respectively, the choice of the government and that of the public as a function of $\mathbf{s}^q$. Note that the public is able condition its choice on $\mathbf{s}^q$ because the signals are disclosed. This will no longer be true in the nondisclosure scenario. Our first, intuitive result is that it is never optimal for the government to distort the screening technology so much that the public never supports implementation:

Lemma 3 *Let $q^{\max}$ solve $\sigma_+^{q^{\max}} = T_P$.*[17] *In equilibrium, $q \in [0, q^{\max}]$.*

In the remainder of this section we will therefore assume with no loss of generality that $q \leq q^{\max}$.

[16] In fact, $\sigma_+^q = \frac{[\theta+q(1-\theta)]^2}{[\theta+q(1-\theta)]^2+[(1-\theta)+q\theta]^2}$, $\sigma^q = \frac{[\theta+q(1-\theta)](1-\theta)}{[\theta+q(1-\theta)](1-\theta)+[(1-\theta)+q\theta]\theta} \leq \frac{1}{2}$ and $\sigma_-^q = \frac{(1-\theta)^2}{V}$. Note that $\lim_{q\to1} \sigma_+^q = \frac{1}{2}$ and $\lim_{q\to1} \sigma^q = 1-\theta$.

[17] It is not hard to show that $q^{\max} = \frac{(1-\theta)-k\theta}{k(1-\theta)-\theta}$, where $k = \sqrt{\frac{1-C_w}{C_w}}$. Clearly $q^{\max} \in [0,1)$ since $C_w \in (\frac{1}{2}, \sigma_+^G]$.

3.1 Strong Electoral Concerns

This subsection studies the game between the government and the public when the electoral costs E are large relative to the private benefits B: $T_{Gov} > \frac{1}{2}$. Since $q \leq q^{\max}$, the public will support implementation if and only if the signals are (α, α). Furthermore, $T_{Gov} > \frac{1}{2}$ implies that the government is disciplined by public opinion.[18] However, as we will see shortly, strong electoral concerns are not sufficient to guarantee that information will be unbiased.

Let $E(\pi_d^q)$ be expected payoff to the government when the size of the bias is q and the government is disciplined by public opinion (hence the subscript d). The government will choose q to maximize $E(\pi_d^q)$ subject to the constraint $q \in [0, q^{\max}]$. However, it is easier to solve the problem in terms of counterfactuals. Specifically, we compare $E(\pi_d^q)$ with $E(\pi_d^G)$, the payoff that the government would get if it was disciplined by public opinion but there was no distortion in information. Define $\Delta\pi_{TH}^q \equiv E(\pi_d^q) - E(\pi_d^G)$, $q \in [0, q^{\max}]$. It can be shown that[19]

$$\begin{aligned} \Delta\pi_{TH}^q &= 2\Pr\left((\alpha,\alpha)^q \mid (\varnothing,\alpha)^G\right)\Pr((\varnothing,\alpha)^G)\left[\pi(w,w \mid (\varnothing,\alpha)^G) - \pi(n,n \mid (\varnothing,\alpha)^G)\right] \\ &\quad - \Pr\left((\alpha,\alpha)^q \mid (\varnothing,\varnothing)^G\right)\Pr((\varnothing,\varnothing)^G)\left[\pi(w,w \mid (\varnothing,\varnothing)^G) - \pi(n,n \mid (\varnothing,\varnothing)^G)\right] \end{aligned}$$

which reduces to

$$\Delta\pi_{TH}^q = qR\left[\frac{1}{2} - (C_w - B)\right] - \frac{1}{2}q^2 V\left(C_w - B - \sigma_-\right). \tag{5}$$

To interpret these equations, suppose that the genuine signals were mixed but $(\alpha,\alpha)^q$ was observed instead (this occurs with probability $2\Pr((\alpha,\alpha)^q \mid (\varnothing,\alpha)^G)\Pr((\varnothing,\alpha)^G) = qR$). The expected payoff to the government in this case is $\pi(w,w \mid (\varnothing,\alpha)^G)$. Note that had the screening technology not being distorted, the expected payoff to the government would have been $\pi(n,n \mid (\varnothing,\alpha)^G)$. Thus the net return from distorting the screening technology in this case is $\pi(w,w \mid (\varnothing,\alpha)^G) - \pi(n,n \mid (\varnothing,\alpha)^G) = \frac{1}{2} - (C_w - B)$. Hence the first term in (5). A similar reasoning applies to all the other signal configurations.

Equation (5) captures the tradeoff that the government faces when the information is

[18] To see this, note that the public and the government always agree on the optimal policy when the evidence is clear-cut. When the evidence is mixed, the government does not implement the project since the public supports n and $\sigma_{Gov} = \Pr(W|(\alpha,\varnothing)^q) \leq \frac{1}{2} < T_{Gov}$ (see Lemma 1).

[19] This and other payoff functions are derived in Appendix A.

distorted. On the one hand, manipulation is costly because the project may be wrongly implemented. On the other hard, the government might want to implement the project when the public does not (this happens when $B > C_w - \frac{1}{2}$ and the genuine signals are mixed). Here a bias in information can help the government: if (α, α) is observed, in fact, the project can be implemented with the support of the public. Note also that when private benefits are small ($B \leq C_w - \frac{1}{2}$), this incentive is absent and the government does not manipulate the screening technology: $\Delta\pi^q_{TH} < 0$.

Let q^{TH} denote the optimum size of the bias. We have

Proposition 4 *If $T_{Gov} > \frac{1}{2}$, the government is always disciplined by public opinion. Furthermore*

$$q^{TH} = \begin{cases} \min\left\{r\frac{1/2-(C_w-B)}{C_w-B-\sigma_-}, q^{\max}\right\} > 0 & \text{if } C_w - B < \frac{1}{2} \\ 0 & \text{if } C_w - B \geq \frac{1}{2} \end{cases}$$

Thus, when B is 'small', citizens enjoy both the benefits of discipline and unbiased information. More generally, we can compute citizen welfare in terms of counterfactuals, exactly as we did for the government. Let $E(U^q_d)$ be the public's expected payoff when the government is disciplined by public opinion and the size of the bias in information is q. We have

$$\Delta U^q_{d,d} \equiv E(U^q_d) - E(U^G_d) = -qR\left(C_w - \frac{1}{2}\right) - \frac{1}{2}q^2V\left(C_w - \sigma_-\right), \tag{6}$$

One can heuristically obtain (6) as follows. Note that biased information affects the public only when the genuine signals are $(\varnothing, \varnothing)$ or mixed but (α, α) is observed instead. In those cases in fact the government policy switches from n to w. To obtain (6) it then suffices to note that these events occur with probability $\frac{1}{2}q^2V$ and qR and yield a welfare loss of $C_w - \sigma_-$ and $C_w - \frac{1}{2}$, respectively. Unsurprisingly, citizen welfare is maximized at $q = 0$, that is, when information is unbiased.

3.2 Weak Electoral Concerns

When electoral concerns are small relative to the private benefits ($T_{Gov} \leq \frac{1}{2}$), two cases can arise. In both situations the government implements the project when the signals are (α, α) and picks the status quo when $(\varnothing, \varnothing)$. The distinction arises when the evidence is mixed. If q is 'small' ($\sigma^q \geq T_{Gov}$), the government feels relatively confident that the α-signal is genuine and implements the project. The opposite is true when q is 'big' ($\sigma^q < T_{Gov}$), and the government is disciplined by public opinion. Note the tradeoff here between discipline and information: As the screening technology becomes more biased, the government changes its decision rule so that its policy is always supported by the public. Interestingly, therefore, biased information may help discipline the government.

The equilibrium of this game can be described formally as follows. Let $\hat{q}$ solve $\sigma^q = T_{Gov}$ (this requires $\frac{1}{2} \geq T_{Gov} > 1 - \theta$). One can show that

$$\hat{q} = r\frac{1/2 - T_{Gov}}{T_{Gov} - \sigma_-}, \quad \text{where} \quad r \equiv \frac{R}{V} = \frac{2\theta(1-\theta)}{\theta^2 + (1-\theta)^2} < 1$$

(see Appendix B). If $T_{Gov} \leq 1 - \theta$, let $\hat{q} = 1$. Note that if the government is not very disposed toward the project ($T_{Gov} \simeq \frac{1}{2}$), then even a small bias in information is conducive to discipline ($\hat{q} \simeq 0$).

Let $E(\pi^q_{nd})$ be the expected payoff to the government when the size of the bias is q and the government is not disciplined by public opinion. For $T_{Gov} \leq \frac{1}{2}$, define $\Delta\pi^q_{TL} : [0, q^{\max}] \to \mathbb{R}$ as follows[20]

$$\Delta\pi^q_{TL} = \begin{cases} \Delta\pi^q_{nd,nd} = E(\pi^q_{nd}) - E(\pi^G_{nd}) & \text{if } q \in [0, \min\{\hat{q}, q^{\max}\}] \\ \Delta\pi^q_{d,nd} = E(\pi^q_{d}) - E(\pi^G_{nd}) & \text{if } q \in (\hat{q}, q^{\max}] \text{ and } \hat{q} < q^{\max} \end{cases} \tag{7}$$

Clearly the government will distort the screening technology whenever $\Delta\pi^q_{TL} > 0$ for some q since then $E(\pi^q_{nd}) > E(\pi^G_{nd})$. Furthermore, any q that maximizes $\Delta\pi^q_{TL}$ will also maximize the government's expected payoff $E(\pi^q_{nd})$. $\Delta\pi^q_{TL}$ is explicitly derived in Appendix A, here we only highlight its main qualitative features.

Proposition 5 *$\Delta\pi^q_{nd,nd}$ is strictly convex in q. $\Delta\pi^q_{d,nd}$ is strictly concave in q and achieves its maximum at $q^* = r\frac{1/2-(C_w-B)}{C_w-B-\sigma_-} \geq \hat{q}$. $\Delta\pi^q_{TL}$ is continuous on $[0, q^{\max}]$ and $\frac{\partial\Delta\pi^q_{d,nd}}{\partial q} > 0$ at $q = \hat{q} < 1$.*

Proposition 5 illustrates an interesting nonconcavity in the government's problem. It implies that the optimum size of the bias will be either zero or $\min\{q^*, q^{\max}\} > 0$, depending

[20] Of course, the second region in (7) does not exist if $\hat{q} \geq q^{\max}$.

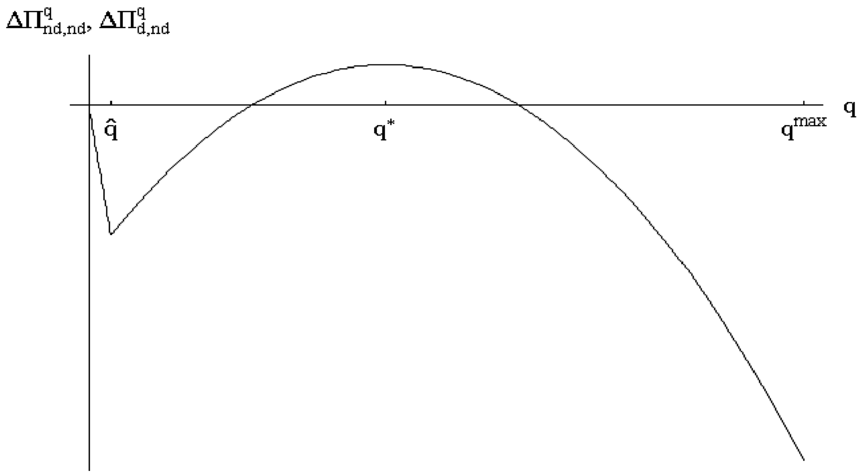

Figure 2: $\Delta\pi^q_{nd,nd}$ and $\Delta\pi^q_{d,nd}$. Parameter values: $\theta = 4/5$, $C_w = 0.6$, $B = 0.27$, $E = 0.15$ ($\hat{q} \simeq 0.02$, $q^* \simeq 0.29$ and $q^{\max} \simeq 0.71$)

on parameter values.[21] Note also that if $q = 0$, then the government is not disciplined by public opinion. Thus, when the electoral concerns are small relative to the private benefits ($T_{Gov} \leq \frac{1}{2}$), the best outcome for the public–no bias in information and discipline–is not achievable.

3.3 Can the Public Benefit from Biased Information?

We can now evaluate citizen welfare when electoral concerns are weak ($T_{Gov} \leq \frac{1}{2}$). As for the government two cases must be considered, depending on whether $q \lesseqgtr \hat{q}$. Let $E(U^q_d)$ $(E(U^q_{nd}))$ be the public's expected payoff when the government is (not) disciplined by public opinion and the size of the bias in information is q. Define $\Delta U^q_{i,j} = E(U^q_i) - E(U^G_j)$, $i, j = d, nd$. We obtain

$$\Delta U^q_{nd,nd} = -(1-q)\, qV\,(C_w - \sigma_-) - \frac{1}{2}q^2 V\,(C_w - \sigma_-) \tag{8}$$

$$\Delta U^q_{d,nd} = (1-q)\, R\left(C_w - \frac{1}{2}\right) - \frac{1}{2}q^2 V\,(C_w - \sigma_-)\,. \tag{9}$$

[21]It is not hard to show that $\Delta\pi^q_{nd,nd}, \Delta\pi^q_{d,nd} \leq 0$ when $E = 0$ (see equations (16) and (17) in the appendix). Thus, we can only have $q > 0$ if E is big enough. For instance, in the example of Figure 2 it suffices to take $E = 0.1$ for the optimal q to be zero.

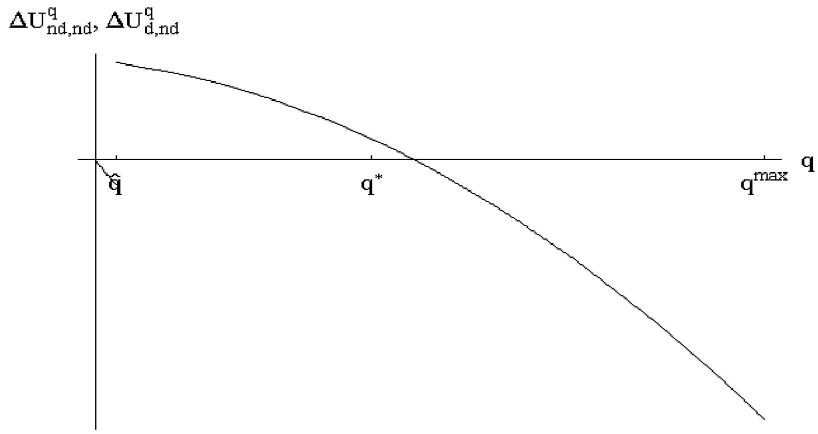

Figure 3: $\Delta U^q_{nd,nd}$ and $\Delta U^q_{d,nd}$ for the same parameter values as in Figure 2

(8) measures citizen welfare when $q < \hat{q}$, and (9) applies when $q > \hat{q}$. In both cases $E(U^G_{nd})$ is the counterfactual since when electoral concerns are weak and $q = 0$ the government is not disciplined by public opinion. We stress that (9) has both positive and negative components. As always, in fact, manipulations decrease citizen welfare as the wrong decision is more likely to be taken. On the other hand, when both the observed *and* the actual signals are mixed (which occurs with probability $(1-q)R$), discipline benefits the public as the government policy switches from w to n. The size of this benefit is given by $C_w - \frac{1}{2}$, as captures by the first term in (9).

The qualitative features of $\Delta U^q_{d,nd}$ and $\Delta U^q_{nd,nd}$ are shown in Figure 3. Both functions are decreasing in q since, for any *given* decision rule, biased information hurts the public. However, when the screening technology is so biased that the government changes its decision rule, the public's payoff 'jumps up'. To see that, write

$$\Delta U^q_{d,nd} - \Delta U^q_{nd,nd} = (1-q)\left[R\left(C_w - \frac{1}{2}\right) + qV\left(C_w - \sigma_-\right)\right] \geq 0. \tag{10}$$

This function measures the value of discipline for given q. Particularly interesting is the case when the size of the bias is small, since then $\Delta U^q_{d,nd} - \Delta U^q_{nd,nd} \simeq R\left(C_w - \frac{1}{2}\right) > 0$, which is just the probability that genuine signals are mixed (R) times the gain from switching from w to n, $C_w - \frac{1}{2}$. Note that since the value of discipline is strictly positive at $q \simeq 0$, if citizens could achieve discipline with little distortion, they would prefer such situation to one in which no manipulations are possible. Figure 2 and 3 illustrate this possibility. Clearly citizens may prefer to tolerate some bias in information (optimally chosen by the government) in order to more closely align the government's interests with its own.[22]

Proposition 6 *In a disclosure environment, a commitment* not *to manipulate the information by the government may hurt the public.*

Proposition 6 shows that information manipulation may have a bright side. The intuition is that, as information becomes more biased toward the government's initial disposition, the government becomes increasingly concerned about forgeries and tends to implement the project less often. As the benefits of discipline can outweigh the losses caused by poor information, paradoxically the public could be worse off if the government committed not to manipulate the information.

4 Full Disclosure with Unobservable Bias

In this section we consider a scenario where signals are disclosed but the public ignores the extent to which the screening technology has been distorted.

We begin with the case when $T_{Gov} > \frac{1}{2}$ and suppose for the moment that the public supports w when (α, α) is observed. Thus the government faces the same problem as in Section 3, with the important difference that now $q \in [0,1]$.[23] Since the government's payoff is given by $\Delta\tilde{\pi}^q_{TH} = \Delta\pi^q_{TH}$ for *all* $q \in [0,1]$, Proposition 4 implies that the optimal size of the bias is 0 if $C_w - B \geq \frac{1}{2}$ and $\min\{q^*, 1\}$ if $C_w - B < \frac{1}{2}$.

Similarly, if $T_{Gov} \leq \frac{1}{2}$ and assuming for now that the public supports w when (α, α) is observed, we can define $\Delta\tilde{\pi}^q_{TL} : [0,1] \to \mathbb{R}$ as

$$\Delta\tilde{\pi}^q_{TL} = \begin{cases} \Delta\pi^q_{nd,nd} & \text{if } q \in [0, \hat{q}] \\ \Delta\pi^q_{d,nd} & \text{if } q \in (\hat{q}, 1] \end{cases} \tag{11}$$

[22] More examples can be constructed as follows. Let $T_{Gov} \simeq \frac{1}{2}$, so that $\hat{q} \simeq 0$. Then $q^* = \arg\max_q \Delta\pi^q_{d,nd} \simeq r\,\frac{E}{C_w - B - \sigma_-}$. Thus if E is small and/or signals are precise, q^*will also be small. Hence we can have $q^* < q^{\max}$ and $\Delta U^{q^*}_{d,nd} > 0$. Finally, note that since $\Delta\pi^{q^*}_{d,nd} \simeq \frac{1}{2}q^* RE > 0$, q^* yields a larger expected payoff to the government than $q = 0$.

[23] In Section 3, by contrast, we restricted attention to $q \in [0, q^{\max}]$ with no loss of generality by Lemma 3.

which is the same as (7), provided $q^{\max} = 1$. Thus it follows from Proposition 5 that the optimal size of the bias will then be either 0 or $\min\{q^*, 1\}$, depending on parameter values.

In both cases, the key issue is whether it makes sense for the public to support w when the signals are (α, α). For instance, if $T_{Gov} > \frac{1}{2}$, the public should expect $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}^q_{TH}$ to be played in equilibrium. It should therefore support w with probability one after observing (α, α) only if $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}^q_{TH} \leq q^{\max}$. In general

Proposition 7 *Suppose* $T_{Gov} > \frac{1}{2}$. *A pure strategy Nash equilibrium (PSNE) of the full disclosure model with unobservable bias exists if and only if* $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}^q_{TH} \leq q^{\max}$.[24] *If* $T_{Gov} \leq \frac{1}{2}$, *a PSNE exists if and only if* $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}^q_{TL} \leq q^{\max}$.

Thus, when the conditions in Proposition 7 are met,[25] the analysis of the previous section goes through unchanged. This is the case in particular for the example in Figure 2 and 3, thus showing that Proposition 6 holds regardless of whether q is observable or not. Matters are more complicated when pure strategy Nash equilibria do not exist, as several cases must be distinguished. Nevertheless, equilibrium strategies share two key features, as described below.

Proposition 8 *Suppose pure strategy Nash equilibria do not exist in the full disclosure model with unobservable bias. A mixed strategy equilibrium exists in which:*

i) The government randomizes among the q*'s belonging to a set* Q *so that in equilibrium the public is indifferent between supporting* w *or* n *when* (α, α) *is observed. The government observes the realization of* q *and then sets the policy optimally.*

ii) Provided (α, α) *is observed, the public randomizes between supporting* w *or* n *so that the* q*'s in* Q *maximize the government's payoff.*

A complete characterization of equilibrium play is provided in the appendix. Nevertheless, it should be clear that the equilibria with observable and unobservable bias share similar qualitative features. In both cases, the government has an incentive to manipulate the

[24] In the nongeneric case where more than one q maximizes $\Delta\tilde{\pi}^q_{TH}$, we take the smallest one.

[25] These conditions are more likely to be fulfilled when the preferences of the public and those of the government are not too dissimilar. Note in fact that if the public is malleable (C_w close to $\frac{1}{2}$), $q^{\max}$ is close to one. Furthermore, the government has also less of an incentive to manipulate the information when B is small since $r\frac{1/2-(C_w-B)}{C_w-B-\sigma_-}$ is increasing in B.

information but such tendency is constrained by the ability of the public to discount biased information. In equilibrium in fact the public's posterior beliefs must be such that, upon observing (α, α), $\sigma_P \geq C_w$, regardless of whether q is observable or not. Intuitively, the government wants the public to support implementation when both signals are favorable. It may make such signals more likely, but must preserve a minimum of trust in official communications for the public to pay any attention to them.

5 Nondisclosure Scenario

We now consider a scenario when the government commits not to disclose any information and q is unobservable. The latter assumption reflects the idea that it must be difficult if not impossible to assess the reliability of the government's sources (i.e., the size of the bias) when no information is released. However, the government policy is observable and this information can be used to update beliefs.

To characterize the equilibrium in this case, more notation is needed. Let $(q, \gamma(\mathbf{s}^q))$ be a strategy for the government, where $\gamma(\mathbf{s}^q)$ is a probability distribution over policies, conditional on $\mathbf{s}^q$. Let β denote a mixed strategy for the public. Suppose that after observing policy $p \in \{w, n\}$ the public believes the government is playing strategy $(\tilde{q}^p, \tilde{\gamma}^p)$. Let $\beta(p, \tilde{q}^p, \tilde{\gamma}^p)$ denote the public's optimal strategy given these beliefs.[26] Loosely speaking, a (perfect Bayesian) equilibrium of this game is a strategy profile $((q^*, \gamma^*), \beta^*)$ and a set of beliefs such that each player's strategy maximizes his payoff given his beliefs and the other player's strategy, and beliefs are formed according to $((q^*, \gamma^*), \beta^*)$ and Bayes' rule (in particular, $(\tilde{q}^p, \tilde{\gamma}^p) = (q^*, \gamma^*)$). If an action is not played in equilibrium, no restriction is imposed on beliefs.

Our first result shows that, so far as the quality of the information is concerned, secrecy is a good thing.

Proposition 9 *In the nondisclosure scenario, assume that after observing p the public holds beliefs $(\tilde{q}^p, \tilde{\gamma}^p)$. Let $\Pi(q)$ be the maximum expected payoff that the government can achieve in that case when the size of the bias is q. Then $\Pi(0) \geq \Pi(q)$.*

[26] The optimal strategy is obtained by first computing the probability that W is the true state of the world (say σ_P) given d, and then choosing w or n depending on whether $\sigma_P \gtrless C_w$ (a nondegenerate mixed strategy is played only if $\sigma_P = C_w$). Clearly, to compute σ_P (via Bayes' rule), the public needs to have some beliefs about (q, γ).

Proposition 9 provides a rationale for restricting attention to equilibria where $q = 0$. It says that for any given belief that the public may hold, knowing more ($q = 0$) cannot hurt the government. Obviously this result relies on all screening technology being equally costly. If the unbiased screening technology ($q = 0$) was prohibitively costly, for instance, the result would not be true. This result should therefore be interpreted as a comparison of gross expected payoffs. Secondly, we emphasize that in games more information is not necessarily beneficial, not even for the player who holds it (see, e.g., Bassan, Scarsini, Zamir (1997)). However, as Kamien, Tauman and Zamir (1990) point out, "it is not the information itself that harms player 1 [the informed player] but the fact that player 2 knew that he had it" (p.133). In the nondisclosure scenario q is not publicly known. Of course, in equilibrium citizens find out q, but for any *given* belief they might hold, the government has the incentive to acquire more information. It is this incentive that drives the result.

We now characterize equilibrium play in the nondisclosure scenario when $q = 0$. Let $\hat{\sigma} = \Pr(W \mid w)$ be the probability that $S = W$ after observing $p = w$ and given that the government picks w with probability one when $\mathbf{s}^G \in \{(\alpha, \alpha), (\alpha, \varnothing), (\varnothing, \alpha)\}$ and n otherwise:

$$\hat{\sigma} = \frac{\theta^2 + R}{1 + R} \in \left(\frac{1}{2}, 1\right). \tag{12}$$

We write $\gamma(\mathbf{s}^q) = p$ if the government selects p with probability one when the signals are $\mathbf{s}^q$. Similarly, we write $\beta(p) = v$ if the public supports v with probability one when the government picks p.

Proposition 10 *The following is a (perfect Bayesian) equilibrium of the nondisclosure game:* $q = 0$, $\gamma((\alpha,\alpha)^G) = w$, $\gamma((\varnothing,\varnothing)^G) = n$, $\beta(n) = n$. *Furthermore,* $\gamma((\alpha,\varnothing)^G) = \gamma((\varnothing,\alpha)^G)$ *and*

i) if $T_{Gov} \leq \frac{1}{2}$ *and* $\hat{\sigma} \geq (<) T_P$, *then* $\gamma((\alpha,\varnothing)^G) = w$ *and* $\beta(w) = w$ (n).

ii) if $T_{Gov} > \frac{1}{2}$, $C_w - B < \frac{1}{2}$ *and* $\hat{\sigma} \geq T_P$, *then* $\gamma((\alpha,\varnothing)^G) = w$ *and* $\beta(w) = w$.

iii) if $T_{Gov} > \frac{1}{2}$, $C_w - B < \frac{1}{2}$ *and* $\hat{\sigma} < T_P$, *then* $\gamma((\alpha,\varnothing)^G) = \begin{cases} w & \text{with probability } \check{s} \\ n & \text{with probability } 1-\check{s} \end{cases}$

and $\beta(w) = \begin{cases} w & \text{with probability } 1-\check{z} \\ n & \text{with probability } \check{z} \end{cases}$ *where* $\check{s} = \frac{\sigma_+^G - C_w}{2r(C_w - 1/2)}$ *and* $\check{z} = \frac{1/2-(C_w-B)}{E}$.[27]

[27] It is not hard to see that $\hat{\sigma} < T_P \Longrightarrow \check{s} < 1$, and that $T_{Gov} > \frac{1}{2}, C_w - B < \frac{1}{2} \Longrightarrow \check{z} < 1$.

iv) if $C_w - B \geq \frac{1}{2}$*, then* $\gamma((\alpha, \varnothing)^G) = n$ *and* $\beta(w) = w$*.*

Not surprisingly, as the government becomes more disposed toward the project, the probability that the status quo is chosen when signals are mixed decreases. For large bias ($T_{Gov} \leq \frac{1}{2}$), this probability is zero and discipline is not attainable. For intermediate bias ($T_{Gov} > \frac{1}{2}$, $C_w - B < \frac{1}{2}$), the government selects the status quo with positive probability if the public is not malleable ($\hat{\sigma} < T_P$) – an outcome that may be termed 'partial discipline'. However, if $\hat{\sigma} \geq T_P$, the government 'gets away' with choosing implementation with probability one. This is because citizens cannot distinguish whether the choice of the government to pick w was due to mixed or two positive signals, and since they are not very biased against implementation, they end up supporting the government. Finally, for small bias ($C_w - B > \frac{1}{2}$), both discipline and no manipulations obtain.

5.1 Transparency vs. Secrecy

Armed with these results, we can compare citizen welfare under disclosure environments. A key implication of the model is that the choice between transparency (full disclosure) and secrecy (nondisclosure) involves a basic tradeoff between the quality of information and discipline. On the one hand, in fact, commitment not to disclose information is very effective at protecting the integrity of the process of information gathering and evaluation. On the other hand, we see that under full disclosure the government is always disciplined by public opinion when the conflict of interest is not large ($T_{Gov} > \frac{1}{2}$) and sometimes even when it is large ($T_{Gov} \leq \frac{1}{2}$).[28] In a nondisclosure environment, by contrast, policymakers are never disciplined when $T_{Gov} \leq \frac{1}{2}$ and may not act according to public desires even for 'intermediate' conflicts of interest ($T_{Gov} > \frac{1}{2}$, $C_w - B < \frac{1}{2}$). Note that such lack of discipline can be traced back to a lack of effective accountability since the public cannot distinguish whether the decision to implement the project was a good one ($\mathbf{s}^G = (\alpha, \alpha)$) or a bad one (e.g., $\mathbf{s}^G = (\alpha, \varnothing)$).

It is not hard to see that either scenario can be optimal. Suppose for instance that $T_{Gov} > \frac{1}{2}$ and $C_w - B < \frac{1}{2}$ (an 'intermediate' conflict of interest). Furthermore, let $\hat{\sigma} \geq T_P$ (C_w small). The public's expected payoff in the full disclosure scenario is $E(U_d^q)$, where q given

[28] See Proposition 8(b) in Appendix B for the case when q is unobservable.

by $q^{TH} = \min\{q^*,\ q^{\max}\} > 0$. $E(U_{nd}^G)$ is the public's expected payoff in the nondisclosure scenario. The difference between these expected payoffs is

$$\Delta U_{d,nd}^{q^{TH}} \equiv E(U_d^{q^{TH}}) - E(U_{nd}^G) = (1-q^{TH})R\left(C_w - \frac{1}{2}\right) - \frac{1}{2}(q^{TH})^2 V\left(C_w - \sigma_-\right).$$

Secrecy performs better than transparency when discipline is not important ($C_w \simeq \frac{1}{2}$) and B is 'large' since then $q^{TH} \simeq r\frac{B}{0.5-B-\sigma_-}$ and $\Delta U_{d,nd}^{q^{TH}} < 0$. By contrast, transparency is preferable when the government is not very biased toward the project ($C_w - B \simeq \frac{1}{2}$) since then $q^{TH} \simeq 0$. Finally, note that when the conflict of interest is small ($C_w - B \geq \frac{1}{2}$), transparency and secrecy both achieve the 'first best' outcome (no distortion and discipline) and are therefore equivalent.

6 Independence of Government Agencies and Optimal Constitutions

A key assumption of the model is that incumbents can easily interfere with the gathering and evaluation of information. This fits the observation that in the US the president appoints and can remove the heads of the executive agencies, thus exerting enormous influence over their policy decisions. Similarly, scholars have argued that "Congress can use appropriations as both carrot and stick—providing additional funds for bureaucrats who produce pleasing decisions and withholding funds for those who do not" (Arnold (1987), p.280). Nevertheless, in reality executive power over government agencies is not without bounds. Of special interest is the growing phenomenon of *independent* government agencies such as the Federal Trade Commission in the U.S. and the Bank of England. Such agencies are not subject to the same degree of political control as other executive agencies and have a status that ensures their independence from political pressure, for instance by limiting the removal of their heads to certain causes.

In this section we focus on the implications of granting *full* independence to the government agencies in charge of collecting and evaluating pre-decision information. Formally, we model the creation of an independent agency as a commitment not to manipulate the information. Thus we posit that if an agency is completely insulated from political pressure, it will carry out its job according to the highest standards of 'objectivity'. This all-or-nothing definition is of course extreme in that 'independence' is often a matter of degree. For instance, the Butler Report provides suggestions on how to strengthen the independence of the Joint Intelligence Committee in the UK but falls well short of recommending its full independence from the executive.[29] With this caveat in mind, it is easy to show that:

Proposition 11 *Suppose the government agencies are independent (i.e., $q = 0$). Then disclosure of information always dominates nondisclosure from the citizens' viewpoint.*

Indeed, if information cannot be manipulated, more information always benefits the public for it makes the government more accountable. Our next step is to find the set of institutional variables (a 'constitution') which maximizes citizen welfare. Table I summarizes the various possibilities.

Table I: Possible Constitutions

I	Full Disclosure & Independent Agencies	II	Full Disclosure & Nonindependent Agencies
III	Nondisclosure & Independent Agencies	IV	Nondisclosure & Nonindependent Agencies

A constitution is said to be optimal if it maximizes citizen welfare. Note that constitution III and IV yield the same welfare since nondisclosure and the independence of government agencies both imply $q = 0$ in equilibrium. Since III is dominated by I, we can thus conclude that disclosure is always a feature of an optimal constitution. Finally, Proposition 6 shows that either constitution I or II can be optimal. We summarize this discussion as follows:

Proposition 12 *An optimal constitution always involves full disclosure of information, but may or may not require independent agencies.*

[29] In matters of national security in fact it is vital to ensure full coordination and cooperation between the executive and the intelligence agencies. Reforms that attempt to insulate intelligence agencies from political pressure may thus not always be feasible or beneficial (see Betts (2004) and the Butler Report, pp. 143-144).

7 Extensions

7.1 Voluntary Disclosure

By definition in the full disclosure model the government shares all its information with the public. One interpretation is that the government *commits* to disclose its signals. However, in reality policymakers often have some discretion as to whether release information. Here we point out that allowing the government to voluntarily disclose some or all of its signals would typically result in all the information being disclosed. Indeed, since information is hard in this model, the logic of Milgrom's (1981) 'unraveling' result applies.

To see this, suppose for simplicity that q is observable and $q \leq q^{\max}$.[30] Beliefs are assumed to be skeptical in the sense that the public expects the government to always disclose favorable information (i.e., an (α, α)-signal), and lack of disclosure is interpreted as evidence that at least one of the signals is $\varnothing$. In this setup, the government has clearly a strict incentive to disclose favorable information. Whether or not the signals are disclosed when $\mathbf{s}^q \neq (\alpha, \alpha)$ is inconsequential since the public will be able to infer that at least one $\varnothing$-signal was observed and hence will support the status quo. As a result, the equilibrium with voluntary disclosure is payoff-equivalent to that of the full disclosure model.

7.2 Domestic vs. Foreign Policy

In an influential book, Shapiro and Page (1992) discuss the possibility that the public might be systematically misled in its policy preferences by government officials, corporations or organized groups. As many of their most prominent examples – including the "Missile Gap" controversy, Tonkin, and the Soviet Scare – involve foreign policy, they argue that

> In matters of foreign policy, the executive branch of government often controls access to information, and it can sometimes conceal or misrepresent reality without being challenged. The political opposition is often intimidated or co-opted. Journalists, even when they are aware of what is going on, sometimes willingly hold back awkward truths in the name of "national security." (p.367).

[30] A similar argument applies to the case when q is unobservable.

In this subsection we briefly outline an extension of the model in which forgeries can be detected, and genuine signals recovered. Consistent with Shapiro and Page's remarks, we also posit that forgeries are more easily detected in domestic policy than in foreign policy. Formally, let ψ_D (ψ_F) be the likelihood that a forged signal concerning domestic (foreign) policy is *not* detected, with $\psi_D < \psi_F$. The probability that a fake signal is actually put into the public domain is thus $\breve{q}_i \equiv q\psi_i$, $i = F, D$. Clearly, by choosing $q \in [0, 1]$ appropriately, the government can control that probability. In particular, using $\breve{q}_i$ as a choice variable, we can go through the same analysis as before, but now with the additional constraint that $\breve{q}_i \leq \psi_i$. Unsurprisigly, allowing for forged information to be detected reduces the likelihood that fake signals are actually put into the public domain. More interestingly, the nonconcavity in the government's program highlighted in Proposition 5 can generate large differences in the equilibrium level of manipulation across policy domains. To see this, consider again the example in Figure 2 and 3. If ψ_i is 'large' (i.e., $\psi_i > q^*$), the equilibrium level of manipulation is the same as in the basic model: q^*. For $\psi_i \leq q^*$, however, the optimal $\breve{q}_i$ gradually decrease with ψ_i until ψ_i reaches approximately 0.16, the value at which $\Delta\pi^q_{d,nd}$ crosses the horizontal axis. After that point, $\breve{q}_i$ suddenly drops to zero. Thus, if $\psi_D < 0.16 < \psi_F$, we may have situations where small differences in the ability to detect forgeries across policy domains generate large differences in equilibrium levels of manipulation.[31]

7.3 More General Biased Technologies

One assumption that we have maintained so far is that only $\varnothing$-signals can be transformed into fake α-signals. This assumption captures the idea that the government will only distort the process of information gathering and evaluation in the direction that support its initial disposition (relative to the public). This subsection discusses how the results of this paper would change should this assumption be dropped. Our answer will hinge crucially upon whether the size of the bias is observable or not.

Consider a scenario where the government appoints a potentially biased expert in charge of providing information. The expert can be biased either toward implementing the project

[31] It should be emphasized in this simple setup the government is not penalized when forgeries are exposed. Relaxing this assumption would clearly reduce the incentive to manipulate the information (thus making transparency more desirable, relative to secrecy). Since this effect would be larger in domestic policy than in foreign affairs (provided $\psi_D < \psi_F$), this subsection probably underestimates the true difference in the incentive to manipulate public opinion across policy domains.

or toward the status quo. In the latter case, let $s \geq 0$ be the probability that the expert transforms a genuine α-signal into fake $\varnothing$-signal. (As before, q denotes the probability that a genuine $\varnothing$-signal is transformed into a fake α-signal by an expert who is biased in favor of the project). The question we ask is whether it can ever be optimal for the government to appoint an expert whose $s > 0$. We begin with the case when s is observable. To see why our answer is affirmative in that case, note that if T_P is very close to $\frac{1}{2}$, then even a small s will induce the public to support implementation when the signal are mixed. As the quality of the signals remain high, the government can thus induce the public to support that decision at little cost. As a result, the government might find it optimal to set $s > 0$.

Things change when the size of the bias is not observable. Suppose in fact that $C_w - B < \frac{1}{2}$ and the public believes that $s > 0$.[32] Can this belief be consistent with the play of the game? The answer is no. Indeed, it is not hard to see that, keeping the public's beliefs fixed, policymakers are better off by setting $s = 0$: fake $\varnothing$-signals can only hurt them.[33] Thus, when the size of the bias is not observable, it is not optimal for the government to distort the screening technology against its initial disposition.

8 Concluding Remarks

This paper develops a simple model where an agent (the government) cares about how its actions are perceived by a principal (the public) and may therefore have an incentive to "fix the evidence" around his a priori favored policy. In such a setting, we studied how important elements of institutional design affect citizen welfare. In line with conventional wisdom, we showed that disclosure of pre-decision information makes the government more accountable and hence more responsive to public desires, relative to a nondisclosure scenario. However, disclosure also induces policymakers to distort the process of information gathering and evaluation, thus compromising the quality of the decision-making process. As a result,

[32]The assumption that $C_w - B < \frac{1}{2}$ is motivated by the fact that if $C_w - B \geq \frac{1}{2}$, then there is always a Pareto-efficient equilibrium in which $s = q = 0$.

[33]To see this, note that the public never supports implementation when the signals are $(\varnothing, \varnothing)$ (since $C_w > \frac{1}{2}$) and always supports implementation when they are (α, α). Furthermore, depending on the public's beliefs about s, the public might either support implementation or the status quo when the signals are mixed. In the former case, since $C_w - B < \frac{1}{2}$, the government achieves its highest possible payoff by just setting $s = q = 0$. In the latter case, setting $s > 0$ is clearly detrimental to the government which might actually choose $q > 0$. In both scenarios the public's beliefs that $s > 0$ are inconsistent with the play of the game.

a commitment to secrecy can sometimes benefit the public. Granting independence from political pressure to the government agencies in charge of collecting and evaluating information always tilts the balance in favor of disclosure. Surprisingly, however, nonindependent agencies can be part of an optimal institutional arrangement. Biased information in fact mitigates the government's tendency to pick its a priori favored policy, and for some parameter values this effect is the dominant one.

The present model is simple and could be extended in several directions. One important extension would be to allow for some information on the consequences of the government policy to be unearthed before citizens make up their minds and protest/vote.[34] Secondly, it could be that even when the decision-making is fully transparent, the government enjoys an informational advantage, relative to the public. For instance, the government may be able to keep a third, informative signal secret (i.e., a confidential/informal report). We conjecture that confidential reports would exacerbate the government's tendency to manipulate publicly available information. However, so long as there is loss of information in manipulating public information, the qualitative results of this paper would clearly still hold. Thirdly, we assumed that different disclosure environments do not affect *per se* the quality of information. However, one of the main virtues of transparency is that independent analysts can verify government information and thinking about an issue. That would not only curb manipulation, but also reduce the risk of 'innocent' mistakes. Finally, and most importantly, the present analysis neglects the important role of bureaucrats' incentives. In the case of performance targets, for instance, there is a clear commonality of interest between the government and individual hospitals in inflating performance figures. This need not always be the case. Studying how both politicians and bureaucrats can be made more accountable and hence responsive to public desire is an exciting direction for future research.

[34] Prat (2005) derives interesting implications along these lines in his model.

Appendix A: Payoffs

In this appendix we derive the expressions in the main text concerning the government's payoff. We omit the derivation of the equations concerning the public's payoff, which can be obtained in a similar fashion.

Recall that $E(\pi^q) = \sum_{\mathbf{s}^q \in \Omega^2} \Pr(\mathbf{s}^q)\pi(p(\mathbf{s}^q), v(\mathbf{s}^q) \mid \mathbf{s}^q)$ and $\Delta\pi^q_{TH} \equiv E(\pi^q_d) - E(\pi^G_d)$, $q \in [0, q^{\max}]$ (the subscript d denotes that the government is disciplined by public opinion). We have

$$\begin{aligned} E(\pi^q_d) &= \frac{1}{2}(\underbrace{[\theta + q(1-\theta)]^2 + [(1-\theta) + q\theta]^2}_{V+2qR+q^2V})\left[-C_w(1-\sigma^q_+) + B\right] \\ &+2\frac{1}{2}(1-q)\underbrace{[(\theta + q(1-\theta))(1-\theta) + ((1-\theta) + q\theta)\theta]}_{R+qV}\left[-C_n\sigma^q\right] \\ &+\frac{1}{2}(1-q)^2V\left[-C_n\sigma_-\right] \end{aligned} \tag{13}$$

since $\Pr((\alpha, \alpha)^q) = \frac{1}{2}([\theta + q(1-\theta)]^2 + [(1-\theta) + q\theta]^2)$, $\Pr((\alpha, \varnothing)^q) = \frac{1}{2}(1-q)[(\theta + q(1-\theta))(1-\theta) + ((1-\theta) + q\theta)\theta]$ and $\Pr((\varnothing, \varnothing)^q) = \frac{1}{2}(1-q)^2V$. Furthermore,

$$\begin{aligned} E(\pi^q_d) &= \frac{1}{2}[V + 2qR + q^2V]\left[-C_w + B\right] + \frac{1}{2}\left[\theta^2 + q^2(1-\theta)^2 + qR\right]C_w \\ &+(1-q)\left[\frac{R}{2} + q(1-\theta)^2\right]\left[-C_n\right] \\ &+\frac{1}{2}(1-q)^2V\left[-C_n\sigma_-\right] \end{aligned} \tag{14}$$

since $\sigma^q_+ = \frac{\theta^2 + q^2(1-\theta)^2 + qR}{V + 2qR + q^2V}$ and $\sigma^q = \frac{\frac{R}{2} + q(1-\theta)^2}{R + qV}$. Now, the first line of (14) can be rewritten as

$$\frac{1}{2}V\underbrace{\left[-C_w + B + \frac{\theta^2}{V}C_w\right]}_{\pi(w,w|(\alpha,\alpha)^G)} + qR\underbrace{\left[-C_w + B + \frac{1}{2}C_w\right]}_{\pi(w,w|(\alpha,\varnothing)^G)} + \frac{1}{2}q^2V\underbrace{\left[-C_w + B + \frac{(1-\theta)^2}{V}C_w\right]}_{\pi(w,w|(\varnothing,\varnothing)^G)},$$

the second line as

$$(1-q)R\underbrace{\left[-\frac{1}{2}C_n\right]}_{\pi(n,n|(\alpha,\varnothing)^G)} + q(1-q)V\underbrace{\left[-C_n\sigma_-\right]}_{\pi(n,n|(\varnothing,\varnothing)^G)}$$

and the third as

$$\frac{1}{2}(1-q)^2 V \underbrace{\left[-C_n\sigma_-\right]}_{\pi(n,n|(\varnothing,\varnothing)^G)}.$$

Thus

$$\begin{aligned} E(\pi_d^q) &= \frac{1}{2}V\pi(w,w|(\alpha,\alpha)^G) + qR\pi(w,w|(\alpha,\varnothing)^G) + \frac{1}{2}q^2V\pi(w,w|(\varnothing,\varnothing)^G) \\ &\quad +(1-q)R\pi(n,n|(\alpha,\varnothing)^G) + q(1-q)V\pi(n,n|(\varnothing,\varnothing)^G) \\ &\quad +\frac{1}{2}(1-q)^2V\pi(n,n|(\varnothing,\varnothing)^G). \end{aligned}$$

It follows that since

$$\begin{aligned} E(\pi_d^G) &= \frac{1}{2}V\pi(w,w|(\alpha,\alpha)^G) + R\pi(n,n|(\alpha,\varnothing)^G) + \frac{1}{2}V\pi(n,n|(\varnothing,\varnothing)^G) \\ &= \frac{1}{2}V\left[-C_w + B + \frac{\theta^2}{V}C_w\right] + R\left[-\frac{1}{2}C_n\right] + \frac{1}{2}V\left[-C_n\sigma_-\right], \end{aligned}$$

$\Delta\pi_{TH}^q \equiv E(\pi_d^q) - E(\pi_d^G)$ is given by

$$\begin{aligned} \Delta\pi_{TH}^q &= qR\left[\pi(w,w|(\alpha,\varnothing)^G) - \pi(n,n|(\alpha,\varnothing)^G)\right] + \frac{1}{2}q^2V\left[\pi(w,w|(\varnothing,\varnothing)^G) - \pi(n,n|(\varnothing,\varnothing)^G)\right] \\ &= qR\left[\frac{1}{2} - (C_w - B)\right] - \frac{1}{2}q^2V\left[C_w - B - \sigma_-\right]. \end{aligned}$$

Now consider the case when $T_{Gov} \leq \frac{1}{2}$. Recall that $\Delta\pi_{TL}^q = \Delta\pi_{nd,nd}^q = E(\pi_{nd}^q) - E(\pi_{nd}^G)$ if $q \in [0, \min\{\hat{q}, q^{\max}\}]$, and $\Delta\pi_{d,nd}^q = E(\pi_d^q) - E(\pi_{nd}^G)$ if $q \in (\hat{q}, q^{\max}]$ and $\hat{q} < q^{\max}$. We have already computed $E(\pi_d^q)$, so we only need to derive $E(\pi_{nd}^q)$. Note that the only difference between $E(\pi_d^q)$ and $E(\pi_{nd}^q)$ is that in the latter case, when the signals are mixed, the government will choose w, not n. The second line of (14) must thus be replaced by

$$\begin{aligned} &(1-q)\left[R+qV\right]\left[-C_w\left(1-\sigma^q\right) + B - E\right] \\ =\ &(1-q)R\underbrace{\left[-C_w + B - E + \frac{1}{2}C_w\right]}_{\pi(w,n|(\alpha,\varnothing)^G)} + (1-q)qV\underbrace{\left[-C_w + B - E + \frac{(1-\theta)^2}{V}C_w\right]}_{\pi(w,n|(\varnothing,\varnothing)^G)} \end{aligned}$$

since $\sigma^q = \frac{\frac{R}{2} + q(1-\theta)^2}{R+qV}$. Thus

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

Using the fact that

$$
\begin{aligned}
E(\pi_{nd}^{G}) &= \frac{1}{2}V\pi(w,w|(\alpha,\alpha)^{G}) + R\pi(w,n|(\alpha,\varnothing)^{G}) + \frac{1}{2}V\pi(n,n|(\varnothing,\varnothing)^{G}) \\
&= \frac{1}{2}V[-C_w + B + C_w\sigma_{+}^{G}] + R[-C_w + B - E + \frac{1}{2}C_w] + \frac{1}{2}V\left[-C_n\sigma_{-}\right]
\end{aligned} \tag{15}
$$

one easily obtains

$$
\Delta\pi_{nd,nd}^{q} = qRE - \frac{1}{2}q^{2}V\left[C_w - B - \sigma_{-}\right] - (1-q)\,qV\left[T_{Gov} - \sigma_{-}\right] \tag{16}
$$

$$
\Delta\pi_{d,nd}^{q} = qRE - \frac{1}{2}q^{2}V\left[C_w - B - \sigma_{-}\right] - (1-q)\,R\left[\frac{1}{2} - T_{Gov}\right]. \tag{17}
$$

Appendix B: Proofs

Proof of Lemma 1. (i) Let $\sigma_{Gov} = \sigma_P = \sigma$. The public supports the decision to implement the project if $\Pr(W|\sigma) \geq C_w$. This implies that $\Pr(W|\sigma) \geq C_w - B + E(1_{\{w,w\}} - 1_{\{n,w\}}) = C_w - B - E$. Part (ii) follows immediately from (2). ■

Proof of Lemma 2. If $v = w$, the claim follows from Lemma 1. $v = n$ implies that $\Pr(W|I) < C_w$ and that the government will choose w if $\Pr(W|\sigma) \geq C_w - B + E$. This can never be the case if $E \geq B$. ■

Proof of Lemma 3. To prove the result, we compare the expected payoff to the government when $q > q^{\max}$ to the payoff it could get by setting $q = 0$. Clearly, if $q > q^{\max}$ the public never supports the decision to implement the project. With strong electoral concerns ($T_{Gov} > \frac{1}{2}$), the government either always chooses the status quo or implements the project if and only if both signals are positive. In both cases, the government's expected payoff is lower than $E(\pi_d^G)$. Note in fact that with $q = 0$ the public would support implementation if both signals were positive and the probability of a mistaken decision would be lower.

Now suppose electoral concerns are weak ($T_{Gov} \leq \frac{1}{2}$). Two cases must be distinguished, depending on whether the government implements the project when the signals are mixed. (We assume that the government implements the project when both signals are positive since scenarios in which the government always picks the status quo are clearly suboptimal.) If the project is implemented when the signal are mixed, we have

$$E(\pi_{nd}^{q>q^{\max}}) =$$

$$\tfrac{1}{2}V\left[-C_w\left(1-\sigma_+^G\right)+B-E\right]+qR\left[-C_w\tfrac{1}{2}+B-E\right]+\tfrac{1}{2}q^2V[-C_w\left(1-\sigma_-\right)+B-E]$$

$$+(1-q)R\left[-C_w\tfrac{1}{2}+B-E\right]+q(1-q)V\left[-C_w\left(1-\sigma_-\right)+B-E\right]+\tfrac{1}{2}(1-q)^2V\left[-C_n\sigma_-\right].$$

Simple algebra then yields

$$E(\pi_{nd}^{q>q^{\max}})-E(\pi_{nd}^{G})=-\frac{1}{2}VE-q\left(1-\frac{q}{2}\right)V\left[C_w-B+E-\sigma_-\right]<0$$

(see (15)). Thus $q>q^{\max}$ cannot be optimal.

If instead the project is not implemented when the signal are mixed, we have

$$E(\pi_{d}^{q>q^{\max}}) =$$

$$\tfrac{1}{2}V\left[-C_w\left(1-\sigma_+^G\right)+B-E\right]+qR\left[-C_w\tfrac{1}{2}+B-E\right]+\tfrac{1}{2}q^2V[-C_w\left(1-\sigma_-\right)+B-E]$$

$$+(1-q)R\left[-C_n\tfrac{1}{2}\right]+q(1-q)V\left[-C_n\sigma_-\right]+\tfrac{1}{2}(1-q)^2V\left[-C_n\sigma_-\right].$$

Thus

$$E(\pi_{d}^{q>q^{\max}})-E(\pi_{nd}^{G})=-\frac{1}{2}VE-(1-q)\,R\left[\frac{1}{2}-T_{Gov}\right]-\frac{1}{2}q^2V\left[C_w-B+E-\sigma_-\right]<0.$$

Again, $q>q^{\max}$ cannot be optimal. ■

Proof of Proposition 5. The government's payoff has been derived in Appendix A. By differentiating $\Delta\pi_{nd}^q$ and $\Delta\pi_d^q$ with respect to q we get $\frac{\partial\Delta\pi_{nd}^q}{\partial q}=RE-qV(C_w-B-\sigma_-)-(1-2q)V\left[T_{Gov}-\sigma_-\right]$, $\frac{\partial^2\Delta\pi_{nd}^q}{\partial q^2}=V(T_{Gov}-\sigma_-+E)>0$, $\frac{\partial\Delta\pi_d^q}{\partial q}=RE-qV(C_w-B-\sigma_-)+R\left[\frac{1}{2}-T_{Gov}\right]$ and $\frac{\partial^2\Delta\pi_d^q}{\partial q^2}=-V(C_w-B-\sigma_-)<0$. Thus $\Delta\pi_{nd}^q$ is strictly convex in q and $\Delta\pi_d^q$ is strictly concave in q. Clearly $q^*\equiv\arg\max_{q\in[0,1]}\Delta\pi_{d,nd}^q=r\frac{1/2-(C_w-B)}{C_w-B-\sigma_-}\geq\hat{q}$.

To show that $\Delta\pi_{TL}^q$ is continuous on $[0,q^{\max}]$, note that $\Delta\pi_{nd}^q$ and $\Delta\pi_d^q$ are both continuous in q on their respective domains. Thus, it suffices to show that $\Delta\pi_{nd}^q=\Delta\pi_d^q$ at $\hat{q}$, $\hat{q}<q^{\max}$.

Recall that $\hat{q}$ solves $\sigma^q = \frac{\frac{R}{2}+q(1-\theta)^2}{R+qV} = T_{Gov}$, provided $T_{Gov} \geq 1-\theta$. (If $T_{Gov} < 1-\theta$, then $\hat{q} = 1$ and the result is trivially true since $\hat{q} \geq q^{\max}$.) This equation yields

$$\hat{q} = \frac{\left(\frac{1}{2} - T_{Gov}\right) R}{VT_{Gov} - (1-\theta)^2} = r \; \frac{1/2 - T_{Gov}}{T_{Gov} - \sigma_-} \tag{18}$$

and one can easily verify that $\Delta\pi_{nd}^{\hat{q}} = \Delta\pi_d^{\hat{q}}$. Simple algebra also yields

$$\left.\frac{\partial \Delta\pi_{nd}^q}{\partial q}\right|_{q=\hat{q}} = RE\left(1 + \frac{\frac{1}{2} - T_{Gov}}{T_{Gov} - \sigma_-}\right) > 0, \quad \hat{q} < 1. \quad \blacksquare$$

Proof of Proposition 7. We prove the result for the case when $T_{Gov} \leq \frac{1}{2}$. The case when $T_{Gov} > \frac{1}{2}$ is similar.

Let $\tilde{q}^{TL} = \arg\max_{q\in[0,1]} \Delta\tilde{\pi}_{TL}^q$ (if there are multiple solutions, take the smallest). Suppose the public believes $q = q^e \leq q^{\max}$. Since $q^e \leq q^{\max}$, the public supports w when (α, α) is observed. The government will therefore maximize $\Delta\tilde{\pi}_{TL}^q$ with respect to $q \in [0,1]$ and choose $\tilde{q}^{TL}$. Consistency requires $q^e = \tilde{q}^{TL}$. Thus we have a PSNE provided $\tilde{q}^{TL} \leq q^{\max}$.

This reasoning fails if $\tilde{q}^{TL} > q^{\max}$. If $q^e \leq q^{\max}$, in fact, it is optimal for the government to pick $\tilde{q}^{TL}$ and beliefs are not consistent with the play of the game. Furthermore, the belief $q^e > q^{\max}$ cannot be supported in equilibrium. This is because the public never supports w if $q^e > q^{\max}$. It is therefore optimal to the government to set $q = 0$ (Lemma 3). Clearly beliefs are again inconsistent with the play of the game since $q = 0$ but $q^e > q^{\max}$. Thus pure strategy Nash equilibria do not exist if $\tilde{q}^{TL} > q^{\max}$. $\blacksquare$

Proposition 8(b) This proposition provides a detailed characterization of the equilibrium sketched in the main text (Proposition 8). Three cases must be distinguished.

Case A. Suppose either $T_{Gov} > \frac{1}{2}$ and $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}_{TH}^q > q^{\max}$ or $T_{Gov} \leq \frac{1}{2}$, $\arg\max_{q\in[0,1]}$ $\Delta\tilde{\pi}_{TL}^q > q^{\max}$ and $\hat{q} < q^{\max}$. The following is a (mixed strategy) equilibrium of the full disclosure game with unobservable bias.

a1) $Q = \{q^{\max}\}$, that is, the government picks $q = q^{\max}$ with probability one. The project is implemented with probability one if $\mathbf{s} = (\alpha, \alpha)$, and is not implemented otherwise.

a2) $$v = \begin{cases} w & \text{with probability } 1 - x^A \text{ if the signals are } (\alpha,\alpha) \text{ and } p = w \\ w & \text{with probability } y \geq x^A \text{ if the signals are } (\alpha,\alpha) \text{ and } p = n \\ n & \text{otherwise} \end{cases}$$

where x^A solves $r\frac{1/2-(C_w-B)-xE}{C_w-B-\sigma_-+xE} = q^{\max}$.

Case B. Suppose $T_{Gov} \leq \frac{1}{2}$, $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}^q_{TL} > q^{\max}$, $\hat{q} \geq q^{\max}$ and $\hat{q} < 1$ ($\Leftrightarrow T_{Gov} > 1-\theta$). Then, in equilibrium

b1) $Q = \{0, q^B\}$, where $q^B \geq \hat{q}$. With probability k^B, the government picks $q = 0$ and then chooses w if $\mathbf{s} \in \{(\alpha,\alpha), (\alpha,\varnothing), (\varnothing,\alpha)\}$ and n otherwise. With probability $1-k^B$, the government picks $q = q^B$ and then choose w if $\mathbf{s} \in \{(\alpha,\alpha), (\alpha,\varnothing), (\varnothing,\alpha)\}$ and n otherwise.

b2) $v = \begin{cases} w & \text{with probability } 1-x^B \text{ if the signals are } (\alpha,\alpha) \text{ and } p = w \\ n & \text{otherwise} \end{cases}$

k^B, q^B and x^B are set as follows. For given $q \in [0,1]$, let $E(\pi^q_{d,x})$ ($E(\pi^q_{nd,x})$) be the expected payoff to the government if the public follows the strategy in (b2) for given x and the government is (not) disciplined by public opinion. x^B is chosen so that $E(\pi^{q^B}_{d,x}) = E(\pi^G_{nd,x})$, where $q^B \equiv \arg\max_{q\in[0,1]} E(\pi^q_{d,x^B})$. k^B is chosen so that $k^B\sigma^G_+ + \left(1-k^B\right)\sigma^{q^B}_+ = C_w$.

Case C. Suppose $T_{Gov} \leq \frac{1}{2}$, $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}^q_{TL} > q^{\max}$ and $\hat{q} = 1$ ($\Leftrightarrow T_{Gov} \leq 1-\theta$). Then, in equilibrium

c1) $Q = \{0,1\}$. $q = 0$ with probability k^C and $q = 1$ with probability $1-k^C$. In both cases, the government chooses w if $\mathbf{s} \in \{(\alpha,\alpha), (\alpha,\varnothing), (\varnothing,\alpha)\}$ and n otherwise.

c2) $v = \begin{cases} w & \text{with probability } x^C \text{ if the signals are } (\alpha,\alpha) \\ n & \text{otherwise} \end{cases}$

x^C and k^C are chosen so that $Q = \arg\max_{q\in[0,1]} E(\pi^q_{nd,x^C})$ and $k^C\sigma^G_+ + \left(1-k^C\right)\sigma^1_+ = C_w$.

Proof of Proposition 8(b). In all three cases, the government's posited strategies are such that in expectation $\Pr(W \mid (\alpha,\alpha)) = C_w$. The public is therefore indifferent between supporting w or n when the signals are (α,α), and the public's posited strategies are optimal. Furthermore, Assumption 2 guarantees that it is optimal for the government to pick $p = n$ when the signals are $(\varnothing,\varnothing)$. The proof is completed by considering each of the three cases in turn.

Case A. Suppose that $T_{Gov} > \frac{1}{2}$ and $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}^q_{TH} > q^{\max}$. For given q, suppose the government is disciplined by public opinion (that is, $p = w$ if and only if (α,α) is observed). Let x be the probability that the public supports n when (α,α) is observed and $p = w$. Following the same steps as in (13), one can show that the government's payoff is given by

$$E(\pi^q_{d,x}) = \frac{1}{2}(V+2qR+q^2V)\left[-C_w(1-\sigma^q_+)+B-xE\right]+(1-q)\left[R+qV\right]\left[-C_n\sigma^q\right] \quad (19)$$
$$+\frac{1}{2}(1-q)^2V\left[-C_n\sigma_-\right]$$
$$= E(\pi^q_d)-\frac{1}{2}(V+2qR+q^2V)xE.$$

Using (5), we get

$$E(\pi^q_{d,x})-E(\pi^G_d)=qR\left[\frac{1}{2}-(C_w-B)\right]-\frac{1}{2}q^2V\left(C_w-B-\sigma_-\right)-\frac{1}{2}(V+2qR+q^2V)xE.$$

Thus, an optimal interior solution q satisfies the first order condition

$$q^*_d(x) \equiv \arg\max_{q\in(0,1)} E(\pi^q_{d,x}) = r\frac{1/2-(C_w-B)-xE}{C_w-B-\sigma_-+xE} \geq \hat{q} \quad (20)$$

since $\hat{q}=r\frac{1/2-(C_w-B)-E}{C_w-B-\sigma_-+E}$ and the second order conditions are always satisfied. Clearly, $q^*_d(x)=q^{\max}$ if we set $x=x^A$ so that

$$q^{\max}=r\frac{1/2-(C_w-B)-x^AE}{C_w-B-\sigma_-+x^AE}. \quad (21)$$

Note that, if $x=0$, by assumption $q^*_d(x)$ is greater than $q^{\max}$. If $x=1$, then the numerator of the above expression is $1/2-(C_w-B)+E=1/2-T_{Gov}<0$. Thus a solution to (21) always exists.

Next, we show that the government wants to implement the project when the signals are $(\alpha,\alpha)^{q^{\max}}$. This is the case because

$$-C_w(1-\sigma^{q^{\max}}_+)+B-x^AE \geq -C_n\sigma^{q^{\max}}_+-yE \quad \Rightarrow \quad B\geq-(y-x^A)E$$

and $y\geq x^A$ (recall that by definition $\sigma^{q^{\max}}_+=T_P=C_w$).

The same equilibrium can be sustained when $T_{Gov}\leq\frac{1}{2}$, $\arg\max_{q\in[0,1]}\Delta\tilde{\pi}^q_{TL}>q^{\max}$ and $\hat{q}<q^{\max}$. Note in fact that since in the posited equilibrium $q=q^{\max}>\hat{q}$ by assumption, the government will be disciplined by public opinion. We therefore have $p=n$ when the signals are mixed.

Case B. Now suppose $T_{Gov}\leq\frac{1}{2}$, $\arg\max_{q\in[0,1]}\Delta\tilde{\pi}^q_{TL}>q^{\max}$, $\hat{q}\geq q^{\max}$ and $\hat{q}<1$. Note that if the government picks $q = q^B$, then it will be disciplined by public opinion (since $q^B \geq \hat{q}$ by assumption). The opposite is of course true if $q = 0$.

The public does not observe the realization of q. Since in equilibrium $\Pr(W \mid (\alpha, \alpha)) = k\sigma_+^G + (1-k)\,\sigma_+^{q^+} = C_w$, citizens are indifferent between supporting w or n if they observe (α, α). Let x be the probability that the public supports n when the signals are (α, α) and $p = w$.

Suppose that, for given q, the government is not disciplined by public opinion. The government's payoff is given by

$$\begin{aligned} E(\pi_{nd,x}^q) &= \frac{1}{2}(V + 2qR + q^2V)\left[-C_w(1-\sigma_+^q) + B - xE\right] + \\ & (1-q)\left[R + qV\right]\left[-C_w\left(1-\sigma^q\right) + B - E\right] + \frac{1}{2}(1-q)^2 V\left[-C_n\sigma_-\right] \end{aligned}$$

and $E(\pi_{nd,x}^q) - E(\pi_{nd}^G)$ is equal to

$$qRE - \frac{1}{2}q^2 V\left[C_w - B - \sigma_-\right] - (1-q)\,qV\left[T_{Gov} - \sigma_-\right] - \frac{1}{2}(V + 2qR + q^2V)xE. \tag{22}$$

By contrast, if the government is disciplined by public opinion, its expected payoff is $E(\pi_{d,x}^q)$ and the optimal choice of q is $q_d^*(x)$ (see equation (19)and (20)). Note that

$$\text{If } \ q < (>)\hat{q}, \ \text{ then } \ E(\pi_{nd,x}^q) \leq (\geq) E(\pi_{d,x}^q). \tag{23}$$

Our goal is to find a x such that (i) $E(\pi_{d,x}^{q_d^*(x)}) = E(\pi_{nd,x}^G)$ and (ii) $E(\pi_{nd,x}^q)$ achieves a maximum at $q = 0$. This implies that choosing $q_d^*(x)$ or 0 gives the government the same expected payoff, given the posited strategies. Consider part (i) first. Note that if $x = 0$, then $E(\pi_{d,0}^{q_d^*(0)}) > E(\pi_{nd,0}^G)$ since the assumptions that $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}_{TL}^q > q^{\max}$ and $\hat{q} < 1$ imply that $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}_{TL}^q \in (\hat{q}, 1]$ (see Proposition 5). If $x = 1$, instead, $E(\pi_{nd,x}^q)$ is strictly decreasing in q. We also have $q_d^*(1) = \hat{q}$ and since $E(\pi_{d,x}^q)$ is strictly decreasing on $q \in [\hat{q}, 1]$, $E(\pi_{d,1}^{q^*(1)}) < E(\pi_{nd,1}^G)$. Thus a x such that $E(\pi_{d,x}^{q^*(x)}) = E(\pi_{nd,x}^G)$ exists. We denote this x by x^B and $q^*(x^B)$ by q^B. Focusing on (ii) now, note that since $E(\pi_{nd,x}^q)$ is convex in q, $E(\pi_{nd,x}^q)$ achieves a maximum either at $q = 0$ or at $q = 1$. By definition of x^B and $q^*(x)$, $E(\pi_{nd,x^B}^G) = E(\pi_{d,x^B}^{q^*(x^B)}) \geq E(\pi_{d,x^B}^1)$. Furthermore, $E(\pi_{d,x^B}^1) \geq E(\pi_{nd,x^B}^1)$ by (23). Therefore, $E(\pi_{nd,x^B}^G) \geq E(\pi_{nd,x^B}^1)$.

Finally, since the government is indifferent between playing 0 and q^B, we can choose k so that $k\sigma_+^G + (1-k)\,\sigma_+^{q^B} = C_w$.

Case C. Now suppose $T_{Gov} \leq \frac{1}{2}$, $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}_{TL}^q > q^{\max}$ and $\hat{q} = 1$. From the strict convexity of $\Delta\tilde{\pi}_{TL}^q = \Delta\pi_{nd,nd}^q$ (since $\hat{q} = 1$), it follows that $\arg\max_{q\in[0,1]} \Delta\tilde{\pi}_{TL}^q = 1$. Thus $E(\pi_{nd,0}^1) > E(\pi_{nd,0}^G)$. Choose x^C so that $E(\pi_{nd,x^C}^1) = E(\pi_{nd,x^C}^G)$. Such x can always be found since $E(\pi_{nd,1}^q)$ is strictly decreasing in q and therefore $E(\pi_{nd,1}^1) < E(\pi_{nd,1}^G)$. By definition of x^C therefore $\{0,1\} \in \arg\max_{q\in[0,1]} E(\pi_{nd,x^C}^q)$. k^C is chosen so that $k^C\sigma_+^G + \left(1-k^C\right)\sigma_+^1 = C_w$.
■

Proof of Proposition 9.[35] Upon observing p, let the public believe that the government is playing $(\tilde{q}^p, \tilde{\gamma}^p)$. Let $\beta(p, \tilde{q}^p, \tilde{\gamma}^p)$ denote the public's optimal mixed strategy, given these beliefs. Suppose the government actually picks $q > 0$. (Since $\tilde{q}^p$ may or may not be equal to q, this need not be an equilibrium.) The best the government can achieve is

$$\begin{aligned}\Pi(q) &= \sum_{\mathbf{s}^q\in\Omega^2} \Pr(\mathbf{s}^q) \max_p \sum_{S\in\Lambda} \pi\left(d, \beta(p,\tilde{q}^p,\tilde{\gamma}^p) \mid S\right) \Pr(S \mid \mathbf{s}^q) \\ &\equiv \sum_{\mathbf{s}^q\in\Omega^2} \Pr(\mathbf{s}^q) \sum_{S\in\Lambda} \pi\left(\gamma^*(\mathbf{s}^q), \beta(p,\tilde{q}^p,\tilde{\gamma}^p) \mid S\right) \Pr(S \mid \mathbf{s}^q).\end{aligned}$$

Note that

$$\Pr(\mathbf{s}^q)\Pr(S \mid \mathbf{s}^q) = \Pr(\mathbf{s}^q \mid S)\Pr(S) = \sum_{\mathbf{s}^G\in\Omega^2} \Pr(\mathbf{s}^q \mid \mathbf{s}^G) \Pr\left(\mathbf{s}^G \mid S\right) \Pr(S) = \sum_{\mathbf{s}^G\in\Omega^2} \Pr(\mathbf{s}^q \mid \mathbf{s}^G) \Pr\left(S \mid \mathbf{s}^G\right) \Pr(\mathbf{s}^G)$$

where the first and third equality follow from Bayes' rule and the second equality follows from $\Pr(\mathbf{s}^q \mid S) = \sum_{\mathbf{s}^G\in\Omega^2} \Pr(\mathbf{s}^q, \mathbf{s}^G \mid S)$ and the 'garbling' condition (3). Thus

$$\Pi(q) = \sum_{\mathbf{s}^G\in\Omega^2} \Pr(\mathbf{s}^G) \sum_{\mathbf{s}^q\in\Omega^2} \Pr(\mathbf{s}^q \mid \mathbf{s}^G) \sum_{S\in\Lambda} \pi\left(\gamma^*(\mathbf{s}^q), \beta(p,\tilde{q}^p,\tilde{\gamma}^p) \mid S\right) \Pr\left(S \mid \mathbf{s}^G\right).$$

Now suppose the government picks $q = 0$. (Since only the policy is observable, the public's beliefs are still given by $(\tilde{q}^p, \tilde{\gamma}^p)$ and the public's best reply is β.) The government's maximum expected payoff is

[35] This is a simple extension of a result in Marschak and Radner (1972), pp. 65-67.

$$\begin{aligned}\Pi(0) &= \sum_{\mathbf{s}^G \in \Omega^2} \Pr(\mathbf{s}^G) \max_p \sum_{S\in\Lambda} \pi\left(d, \beta(p, \tilde{q}^p, \tilde{\gamma}^p) \mid S\right) \Pr(S \mid \mathbf{s}^G) \\ &\equiv \sum_{\mathbf{s}^G \in \Omega^2} \Pr(\mathbf{s}^G) \sum_{S\in\Lambda} \pi\left(\gamma^*(\mathbf{s}^G), \beta(p, \tilde{q}^p, \tilde{\gamma}^p) \mid S\right) \Pr(S \mid \mathbf{s}^G).\end{aligned}$$

From

$$\textstyle\sum_{S\in\Lambda} \pi\left(\gamma^*(\mathbf{s}^G), \beta(p, \tilde{q}^p, \tilde{\gamma}^p) \mid S\right) \Pr(S \mid \mathbf{s}^G) \geq \sum_{S\in\Lambda} \pi\left(\gamma^*(\mathbf{s}^q), \beta(p, \tilde{q}^p, \tilde{\gamma}^p) \mid S\right) \Pr\left(S \mid \mathbf{s}^G\right)$$

and $\sum_{\mathbf{s}^q \in \Omega^2} \Pr(\mathbf{s}^q \mid \mathbf{s}^G) = 1 \; \forall \mathbf{s}^G$, it follows that $\Pi(0) \geq \Pi(q)$. ■

Proof of Proposition 10. By Proposition 9, $q = 0$ maximizes the government's expected payoff for any strategy that the public might play. It is straightforward to check that the remaining strategies form an equilibrium. For the sake of brevity, we will focus on part (iii). Suppose $T_{Gov} > \frac{1}{2}$, $C_w - B < \frac{1}{2}$ and $\hat{\sigma} < T_P$. Clearly $\gamma((\varnothing, \varnothing)^G) = n$ since

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

It is also not hard to see that $\beta(n) = n$. Now note that if $\gamma((\alpha, \varnothing)^G) = w$ with probability one then $\beta(w) = n$ since $\hat{\sigma} < T_P$. But then $\gamma((\alpha, \varnothing)^G) = w$ cannot be optimal since $T_{Gov} > \frac{1}{2}$. Similarly, if $\gamma((\alpha, \varnothing)^G) = n$ with probability one then $\beta(w) = w$ since $\sigma_+^G > T_P$. But then $\gamma((\alpha, \varnothing)^G) = n$ is not optimal since $C_w - B < \frac{1}{2}$. We therefore need mixed strategies. The indifference condition for the government is

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

The indifference condition for the public is $\Pr(W \mid w, \check{s}) = C_w$, where $\Pr(W \mid w, \check{s})$ is the probability that $S = W$ after observing $p = w$ given that the government plays w with probability one if $(\alpha, \alpha)^G$, with probability $\check{s}$ if the signals are mixed and with probability zero if $(\varnothing, \varnothing)^G$. Using Bayes' rule

$$\Pr(W \mid w, \check{s}) = \frac{\theta^2 + \check{s}R}{V + 2\check{s}R} = C_w \implies \check{s} = \frac{\theta^2 - C_w V}{R\,(2C_w - 1)} = \frac{\sigma_+^G - C_w}{2r\,(C_w - 1/2)}.$$

To conclude the proof, note that if the government plays w with positive probability when signals are mixed, it must play w with probability one when $(\alpha, \alpha)^G$. ■

References

[1] Arnold, D. R. (1987). 'Political Control of Administrative Officials', *Journal of Law, Economics and Organization*, vol. 3(2) (Fall), pp. 279-286.

[2] Bassan, B., Scarsini, M. and Zamir, S. (1997). '"I Don't Want to Know": Can It Be Rational?', mimeo.

[3] Besley, T. and Burgess R. (2001). 'Political Agency, Government Responsiveness and the Role of the Media', *European Economic Review*, vol.45(4-6), pp.629-640.

[4] Besley, T. and Burgess R. (2002). 'The Political Economy of Government Responsiveness: Theory and Evidence from India', *Quarterly Journal of Economics*, vol.117(4), pp.1415-1452.

[5] Betts, R.K. (2004). 'The New Politics of Intelligence', *Foreign Affairs*, vol.83(3) (May/June), pp.2-8.

[6] Biglaiser, G. and Mezzetti, C. (1997). 'Politicians' decision making with re-election concerns', *Journal of Public Economics*, vol.66(3), pp.425-447.

[7] Blix, H. (2005). *Disarming Iraq*, Bloomsbury Publishing PLC, London.

[8] BMA survey of A&E waiting times (2005), downloadable at http://www.bma.org.uk /ap.nsf/Content/AandEwaiting.

[9] 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(3) (July), pp.532-550.

[10] 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/.

[11] Cremer, J. (1995). 'Arm's Length Relationships', *Quarterly Journal of Economics*, vol.110(2) (May), pp.275-295.

[12] Dewatripont, M. and Tirole J. (1999). 'Advocates', *Journal of Political Economy*, vol.107(1), pp.1-39.

[13] The Economist (2004). 'Special Report Intelligence Failures', (July 17th) pp.23-25.

[14] The Economist (2005). 'Government Figures', (March 26th), pp.31-32.

[15] Frankel, M. (2001). 'Freedom of Information: Some International Characteristics', Working Paper, The Campaign for Freedom of Information.

[16] Garicano, L. and Posner R. A. (2005). 'Intelligence Failures: An Organizational Economics Perspective', forthcoming *Journal of Economic Perspectives*.

[17] Holmstrom, B. (1979). 'Moral Hazard and Observability', *Bell Journal of Economics*, vol.10 (Spring), pp.74-91.

[18] Holmstrom, B. (1999). 'Managerial Incentive Problems: A Dynamic Perspective', *Review of Economic Studies*, vol.66(1) (January), pp.169-182.

[19] Kamien, M. I., Tauman, Y. and Zamir S. (1990). 'On the Value of Information in Strategic Conflict', *Games and Economic Behavior*, vol.2, pp. 129-153.

[20] Jones, F. and Jones A.-M. (2005). 'Public Confidence in Official Statistics: An analysis based on data collected in the National Statistics Omnibus Survey', downloadable at http://www.statistics.gov.uk/about/data/public_confidence/reports.asp

[21] Levy, G. (2005). 'Decision Making in Committees: Transparency, Reputation and Voting Rules', Working Paper.

[22] Maggi G. and Rodriguez-Clare A. (1995). 'Costly Distortion of Information in Agency Problems', *Rand Journal of Economics*, vol.26(4) (Winter), pp.675-689.

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

[24] Maskin, E. and Tirole J. (2004). 'The Politician and the Judge', *American Economic Review*, vol. 94(4) (September), pp.1034-1051.

[25] Milbourn, T.T., Shockley, R.L. and Thakor, A.V. (2001). 'Managerial Career Concerns and Investments in Information', *Rand Journal of Economics*, vol.32(2), pp.334-351.

[26] Milgrom, P. (1981). 'Good News and Bad News: Representation Theorems and Applications', *Bell Journal of Economics*, vol.12(2) (Autumn), pp.380-391.

[27] Page, B.I. and Robert Y.S. (1992). *The Rational Public: Fifty Years of Trends in Americans' Policy Preferences*, University of Chicago Press, Chicago.

[28] Persson, T., Roland, G. and Tabellini, G. (1997). 'Separation of Powers and Political Accountability', *Quarterly Journal of Economics*, vol.112 (November), pp.1163-1202.

[29] Prat, A. (2005). 'The Wrong Kind of Transparency', *American Economic Review*, vol. 104(3) (June), pp. 946-59.

[30] Prat, A. (2006). 'The More Closely we are Watched, the Better we Behave?' in *Transparency: The Key to Better Governance?*, C. Hood and D. Heald eds., Oxford University Press, Oxford.

[31] Prendergast, C. (1993). 'A Theory of "Yes Men"', *American Economic Review*, vol. 83(4) (September), pp. 757-770.

[32] 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.

[33] 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

[34] Stasavage, D. (2004). 'Open-Door or Closed-Door? Transparency in Domestic and International Bargaining', *International Organization*, vol. 58, pp. 667-703.

[35] Stasavage, D. (2006). 'Does Transparency Make a Difference? The Example of the European Council of Ministers' in *Transparency: The Key to Better Governance?*, C. Hood and D. Heald eds., Oxford University Press, Oxford.

[36] Statistics Commission, Report No.24 (2005). ‘Official Statistics: Perceptions and Trust’, downloadable at http://www.statscom.org.uk/uploads/files/reports/024%20-%20Perceptions%20Trust.pdf

[37] Suurmond, G., Swank O.H. and Visser B. (2004). ‘On the Bad Reputation of Reputational Concerns’, *Journal of Public Economics*, vol.88, pp.2817-2838.

[38] Swank, J., Swank, O. and Visser, B. (2006). ‘Transparency and Pre-meetings’, mimeo.

[39] U.S. Supreme Court (1974). ‘United States v. Nixon, 418 U.S. 683’, downloadable at http://caselaw.lp.findlaw.com/scripts/getcase.pl?navby=CASE&court=US&vol=418&page=683