



Auditing, disclosure, and verification in decentralized decision problems[☆]



Luca Anderlini^{a,*}, Dino Gerardi^b, Roger Lagunoff^a

^a Georgetown University, USA

^b Collegio Carlo Alberto, Italy

ARTICLE INFO

Article history:

Received 5 February 2016

Received in revised form 30 August 2016

Accepted 9 September 2016

Available online 30 September 2016

JEL classification:

C73

D63

D72

D74

H11

Keywords:

Auditing

Disclosure

Agency bias

Ideological bias

ABSTRACT

We study the relative performance of disclosure and auditing in decentralized institutions. We consider the information transmission problem between two decision makers who take actions in sequence at two decision dates. The first decision maker has private information about a state of nature that is relevant for both decisions, and sends a message to the second. The second decision maker can *commit* to *only* observe the *message* (disclosure), or can retain the option to observe the action of the first decision maker (auditing) or, at some cost, to verify the state. In equilibrium, state verification will never occur and the second decision maker effectively chooses between auditing and disclosure.

When the misalignment in preferences reflects an agency bias – a bias in a decision maker's *own action* relative to that of the other – then the second decision maker chooses to audit in equilibrium. When the misalignment in preferences reflects an ideological bias – one decision maker prefers *all actions* to be biased relative to the other decision maker – then, for a large enough misalignment, the second decision maker chooses disclosure in equilibrium. Our results indicate that the ability to commit to forego the audit has value in the latter case.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

In any ongoing organization or society there are two primary ways in which information is revealed. One is *disclosure* – transmission of information by the informed parties to uninformed receivers.¹ Disclosure consists of all forms of manufactured transmission including oral and written (files, archives, etc.) communication. The act of disclosing information may or may not be obligatory.² The key is that its content can be freely chosen by the informed sender.

[☆] A previous draft was circulated under the title: “Do Actions Speak Louder than Words? Auditing, Disclosure, and Verification in Organizations.” We wish to thank Tom Gresik, an Associate Editor, and two anonymous referees for their helpful comments and suggestions. We also thank seminar participants at Johns Hopkins, North Carolina, PUC (Chile), HKUST, the IADB, Universidad Alberto Hurtado, Vanderbilt, and the Midwest Theory Meetings at Indiana University. All three authors are grateful to the National Science Foundation for financial support (Grant SES-0617789). Some of the work was carried out while Luca Anderlini was visiting EIEF in Rome. He gratefully acknowledges their generous hospitality.

* Corresponding author at: Georgetown University, 37th and O Streets NW, Washington, DC 20057, USA.

E-mail address: luca@anderlini.net (L. Anderlini).

¹ Notice that while we use the term disclosure as synonymous with cheap talk, it has also been used in the literature as revelation of information that is *verifiable* by the receiver.

² Record keeping requirements in firms fit this definition.

The other form of revelation is *discretionary auditing* – the attempt by the uninformed party to investigate or verify the information in some fashion. Auditing may take place under the guise of a formal process or institution such as the General Accountability Office in the U.S. government, or it may take the form of an informal monitoring arrangement in a collective action problem.

This paper studies the relative performance of disclosure and auditing in decentralized decision settings. In these settings, it is often decision makers themselves who make choices about which of the two information channels is best.

In firms, for instance, different managers are promoted to key decision-making posts over time. It is standard practice for a current managerial team with inside knowledge to provide a detailed prospectus of the company's overall health periodically. Even so, its policies, expenditures, and acquisitions may also be audited to provide some external verification to shareholders and subsequent generations of managers.

Political systems operate in much the same way. Rival political parties vie for power over time. A ruling political party holds some inside information (for instance, knowledge of the effects of a particular regulation) and is required to disclose what it knows or observes. A rival party eventually reaches power and must then choose to either accept the outgoing party's words or instead investigate the outgoing party's policies and decisions.

Other examples include common pool problems in which users with current access to the resource must decide on the best means to extract information about the stock from earlier users, and regulatory environments in which regulators must decide on how structure reporting requirements.

With these examples in mind, we posit a two-period policy setting model that attempts to capture some of the key trade offs in the choice of transmission mechanism. We assume that the policy at each date is chosen by a distinct decision maker. Disclosure is assumed to take the extreme form of cheap talk, while auditing takes the extreme form of being able to directly observe the actions taken by others. In period 1, the initial decision maker – DM1 (the “sender”) – chooses an action (a “policy”) and a cheap-talk message after privately observing the realization of an underlying state variable that affects the preferences of both DMs. In period 2, the second decision maker – DM2 (the “receiver”) – chooses a policy.

The choices of possible information channels open to DM2 are as follows. She can choose to commit to *only* observe DM1's message *before* DM1 takes any action or sends his message. We refer to this as a commitment to rely on *disclosure*. The decision by DM2 of whether or not to rely on disclosure is observed by DM1 who then chooses his policy. If DM2 chooses not to commit, she retains the option to observe DM1's action, which we refer to as *auditing*, or directly verify the state – referred to simply as *verification*. For simplicity, auditing is assumed to be costless and state verification has a vanishing small (lexicographic) cost.³ Once period 2 arrives, DM2 is in charge of the organization and her action is then taken.

Misaligned preferences between the two temporally separated agents (e.g., two managerial teams, two political parties, or two resource users) distort incentives away from profit-maximizing, or welfare maximizing policies. The effect of this misalignment here is that full disclosure does not occur. In this sense our setup is related to the classic cheap talk environment in Crawford and Sobel (1982) (henceforth “CS”). Their sender is our DM1, while DM2 is their receiver. Unlike in the CS environment, policy decisions here are made by both players and the receiver may choose to audit before taking her own decision.

We ask when disclosure will be used, and when instead organizations rely on auditing, or even verification, in equilibrium. The answer is nuanced and depends on the particular characteristics of the organization.

Our first result, Proposition 1, shows that DM2 will never choose the direct verification option in equilibrium. Intuitively, if DM1 expects the state to be verified directly, he would choose his most preferred first period policy tailored to the state. But because his policy would fully reveal the state, the uninformed DM2 would then switch to the less costly policy audit. This is shown in any pure strategy equilibrium, but we later verify that it holds generically (in the level of bias/misalignment) for mixed strategy equilibria as well.⁴

As a result of Proposition 1, DM2's effective choice is between “words” (disclosure) and the “actions” (policy audit) of her predecessor, DM1. The choice between words and actions ultimately depends on the wedge in preferences between the two parties. We refer to this misalignment as the “bias” (following CS) and examine the following two canonical cases.

In the case of *ideological bias*, the first decision maker DM1 always prefers systematically higher policies than DM2, regardless of which party chooses or which period the choice is made. The ideological biases arise naturally in politics where competition between ideologically differentiated parties takes place.

Ideological bias contrasts with the case of *agency bias* whereby the active decision maker prefers a higher (or lower) policy than his counterpart, the passive stakeholder. This means that DM1 (DM2) prefers, say, a higher action in period 1 (period 2) than DM2 (DM1). Arguably, many private firms conform to this case since the future decision maker (DM2) is often selected internally from existing management, and the previous manager holds stock options that incentivizes his future decisions.

These two forms of bias give rise to quite different transmission protocols. With agency bias, DM2 will generally choose a policy audit. However, under large enough ideological bias, DM2 will commit in advance to use *only* the disclosed information. In other words, we find that under ideological bias, the uninformed party will rely solely on the informed party's cheap talk in equilibrium and thus forego all subsequent audit and state verification opportunities.

³ The assumption of costless policy audits is shown to be without loss of generality. All the results go through if audits costs are small. Our main result (Proposition 3(i)) goes through with arbitrary auditing costs. See Section 5 for a detailed discussion.

⁴ See Section 5.

To gain some intuition for these results, we observe first that equilibria in which DM2 relies only on messages will resemble those in standard cheap talk models (e.g. CS). That is, while DM1 can, in principle, report false information, she chooses not to in equilibrium. Instead, the sender's message is partially, but not fully, informative. The lack of full disclosure is of course bad for DM2. In contrast, when a policy audit occurs in equilibrium, DM1 chooses his action taking into account both its direct payoff effect, driven by the bias, and its signaling value driven by DM1's preferences over DM2's action in the second period.

When the wedge between preferences takes the form of an agency bias, these two effects roughly offset one another, thereby mitigating the effect of the sender's bias. However, in the case of ideological bias, the direct payoff effect and the signaling value reinforce one another; both induce actions that are too large for the recipient. If the reinforced effect of the bias is large enough, the audit becomes undesirable for DM2 who then avoids it in equilibrium. In fact, we show that forgoing the auditing and direct verification options in favor of disclosure is not just preferable for the recipient under ideological bias, but it is Pareto improving for all parties if the ideological bias is large enough.

A critical part of the argument is that DM2 can commit in advance to forego or at least limit the scope of a discretionary audit. To some extent, this can be viewed as a simplifying assumption since the commitment to forego auditing is shown to be preferred by both decision makers in most cases. The commitment becomes institutionalized when a consensus exists among all the parties *ex ante*. Standard corporate policy, for instance, places auditing authority in the hands of the firm's board of directors, rather than the CEO.⁵ Governments make allowances as well to prevent public audits for reasons of national security.

Naturally, in some cases commitment is not always possible. The results of the paper inform these cases as well: without commitment, discretionary auditing will occur making all parties worse off under a large enough ideological bias. At this point the loss of welfare relative to the disclosure regime may well become apparent to the participants, but it simply is too late to avoid it.

In our view this is at least part of the reason why "audits" in the political sphere are often viewed as "partisan witch hunts" where "all actors involved in accountability processes use a variety of strategies to argue their case and apportion blame" (Boin et al., 2008). A "witch hunt" of this type can be used to uncover embarrassing facts about one's political rival. In the language of the model, the more precise is the embarrassing information, the easier it is for the investigating party to implement its ideal policy choice.

Indeed, the frequency of partisan investigations is an indication that institutional commitments to avoid auditing/investigating are more difficult in political systems.

"Partisan investigations have historically been a pox on both houses [in the U.S. Congress] – embarrassing the investigators as much as the investigated and wounding the majority party. [...] Too often, investigative hearings turn out to be flights of vanity for a committee's members to speechify and humiliate witnesses." – Goldberg (2010)

The destructive nature of such investigations suggests a paradox. In firms (where agency bias is prevalent), commitment to forgo auditing may be possible though unnecessary since audits will always be chosen in equilibrium regardless of whether commitment is possible. In political systems, audits often assume the form of a partisan investigation. In this case the commitment to forgo auditing would be chosen if feasible and would in fact be preferable for all parties if the bias is large enough. Unfortunately, commitments of this type are not usually codified in political systems and may thus be infeasible.

The remainder of the paper is organized as follows. Section 2 discusses related literature. Section 3 sets out the model. In Section 4 we characterize equilibria under both ideological bias and agency bias. Section 5 examines the robustness of the main results, including the extension to mixed strategies. Section 6 concludes with a brief discussion of the accountability systems found in firms, polities, and informal collective action scenarios. All proofs are relegated to an Appendix to the paper. A prefix of "A" in the numbering of equations and so on indicates that the relevant item can be found in Appendix.

2. Related literature

We examine the equilibrium choice of disclosure, auditing, or direct verification by the uninformed agent. With some exceptions (discussed below) most of the literature examines one of these transmission mechanisms, but do not avail all the options to the modeled agents.

There is a well established literature on (costly) state verification beginning with Townsend (1979) (see Bolton and Dewatripont, 2005 for further references). By way of comparison, the relevant feature in our model is that state verification does *not* arise in equilibrium when auditing and disclosure are available, even when verification costs are small.

Disclosure is the main focus of the cheap talk communication literature beginning with CS.⁶ The present paper contains elements of dynastic cheap talk reminiscent of Spector (2000) and Anderlini et al. (2012) (AGL) which examine the effects of bias among multiple decision makers and the incentives to disclose the contents of one's objective signals (via cheap talk) to others in the temporal chain. In these papers, cheap talk (disclosure) was the only option for transmitting information; neither auditing or direct verification were available to the uninformed agents.

⁵ See [13].

⁶ See Farrell and Rabin (1996), Sobel (2007), and Krishna and Morgan (2008) for surveys and references.

Conditionally on DM2 *not* committing to disclosure, the present model is connected to the vast literature on signaling started by Spence (1973), since auditing involves the verification of payoff relevant decisions. A model related to ours in the signaling genre comes from Carrillo and Mariotti (2000) who examine a single agent problem with present-biased preferences. Given the bias, their model may be re-interpreted as a multiple agent decision problem – as in the present model. Their focus is on “voluntary ignorance,” i.e., when/whether an agent chooses to stop learning. In a sense, the present model also concerns voluntary ignorance – in this case by the receiver who might choose to commit to disclosure and hence not to “learn” from the sender. Daughety and Reinganum (2010) focus on a related question from a social welfare perspective. They examine a planner’s choice of protocol in a public goods provision problem between privacy rights of an individual and the social benefit of exposing free riders. Our “words versus actions” trade off somewhat resembles their “privacy versus public information” trade-off even though there is no (cheap talk) communication in the latter.

A number of papers have compared outcomes when different types of observation technologies are varied. Prat (2005) compares different models in which a principal can observe either the action directly or instead a noisy signal. Other papers have examined different types of communication protocols that may be chosen by the sender. Austen-Smith and Banks (2000) and Kartik (2007) allow the sender to choose among an array of both costless and costly messages, the latter referred to as “money burning.” Their results demonstrate how the precision of cheap talk increases with the addition of the money burning option. Karamychev and Visser (2011) go a step further in giving a full characterization of the optimal ex ante equilibrium for the sender in the model with both cheap talk and money burning. Kartik (2009) incorporates lying costs into the cheap talk model thereby turning cheap talk into costly signaling. He shows that equilibria exist exhibiting full separation in certain regions of the type space, something that is not possible in the standard CS model with a continuum of types.

Our paper differs from these in that we examine the signaling value of policies rather than that of pure money burning. In this sense, the present paper is more related to Daughety and Reinganum (2008) who study the endogenous choice of protocol by firms that attempt to reveal quality of their products. In their model, a firm can disclose quality through direct claims, or it can signal quality through its product choices.

The main difference between these papers and ours is that we place the choice of communication protocol in the hands of the receiver rather than the sender or a planner. While protocol decisions by the sender are quite natural in a market settings – firms vis a vis consumers – studied by Daughety and Reinganum (2008) and Austen-Smith and Banks (2000) among others, we think it quite natural that the reverse would be likely *within* firms and governments where dynastic considerations play a role. In these cases, the actors fear their policies may be undone by future decision makers if the information were revealed. Hence their actions would not ordinarily come to light unless an explicit audit makes it the case.

3. The model

There are two Decision Makers – DM1 and DM2. The first has decision authority at time $t = 1$ and the second at time $t = 2$. The actions chosen by the two DMs in the two time periods are denoted by $a_1 \in \mathbb{R}$ and $a_2 \in \mathbb{R}$ respectively.

3.1. Payoffs

The symbol $\theta \in \mathbb{R}$ denotes the value of a state of nature that is drawn once and for all from a strictly positive continuous density f defined on the real line \mathbb{R} . The state θ is drawn before anything else takes place.

The two decision makers are differentiated by a bias parameter $b > 0$. This parameter is the source of the preference misalignment between the two DMs. Hence, it is the key driver of all equilibrium decisions in the model. We consider two canonical cases.

The first we call *Ideological Bias*, with preferences of DM1 and DM2, respectively, given by

$$V_1 = -(a_1 - \theta - b)^2 - (a_2 - \theta - b)^2, \quad \text{and} \quad V_2 = -(a_1 - \theta)^2 - (a_2 - \theta)^2 \quad (1)$$

Here, the ideal policy of DM1 in each period $t = 1, 2$ is $a_t = \theta + b$ whereas the ideal policy for DM2 is $a_t = \theta$. We refer to DM1 as being “biased” in both periods since his ideal policy is “distorted” upward by b .⁷

We term the second canonical case *Agency Bias*, with preferences of DM1 and DM2, respectively, given by

$$V_1 = -(a_1 - \theta - b)^2 - (a_2 - \theta)^2, \quad \text{and} \quad V_2 = -(a_1 - \theta)^2 - (a_2 - \theta - b)^2 \quad (2)$$

Under agency bias, the preference misalignment is driven by the role that each DM has in each period: active policy maker or passive stakeholder. Each active DM is biased for the action that corresponds to the period in which he is the active policy maker, and unbiased when he is a “passive stakeholder.” Hence, DM t ideally prefers policy $a_t = \theta + b$ in each period $t = 1, 2$, whereas DM τ ($\tau \neq t$) ideally prefers $a_t = \theta$. Our terminology in this case is motivated by the fact that agency problems arise from the misalignment of preferences between someone who actually controls a payoff-relevant variable (the agent, or the

⁷ Of course, the “bias” label is relative. Moreover, the actual *direction* of the preference misalignment – DM1’s ideal point is larger than DM2’s – is inessential. We could allow for $b < 0$.

active DM in each period in our model) and a passive stakeholder who is affected by the choice but has no direct control over it (the principal, or, in our model, DM2 in period 1 and DM1 in period 2).

Our terminology is intended to capture inherent incentive conflicts in various collective decision settings. As a stylized representation, ideological bias can be found in political institutions in which an outgoing political party has systematically different policy preferences than its newly elected rival. Agency bias can be found in firms in which a manager's preferences are misaligned with his or her successor or predecessor, and DM2 is selected internally from existing management. This new manager, moreover, will have preferences over the use of managerial perks (say private jets), while a past manager who holds stock options awarded during his tenure as a manager only has preferences for managerial actions that induce stock-appreciation. The agency bias case can also apply to politicians provided they are primarily *office-motivated* rather than ideologically driven. Both types bias can be found in informal groups that coalesce to solve collection action problems.⁸

The biases of the the two DMs in the two periods in the two canonical cases we have identified are represented schematically below.

	$\tau = 1$	$\tau = 2$		$\tau = 1$	$\tau = 2$
DM1	b	b	DM1	b	0
DM2	0	0	DM2	0	b
Ideological Bias			Agency Bias		

(3)

Note that in either the ideological or the agency bias case, as $b \rightarrow 0$ the conflict between the players vanishes.

3.2. Information, auditing, and disclosure

In this subsection, we introduce the formal auditing model. The focus is on how the degree of conflict – as represented here by b – distorts the incentives of decision makers at each decision date.

Crucially, we assume that neither the state θ nor the first period policy is automatically observed by DM2. Specifically, the state θ and the first period policy a_1 are known only to DM1, unless DM2 makes it a point to verify them. DM2 therefore assumes power in $t = 2$ knowing only that θ is realized from a continuous strictly positive density f defined on \mathbb{R} .

After observing θ and choosing his policy a_1 , the informed agent DM1 chooses a message $m \in \mathbb{R}$ that can in principle be used to communicate or “disclose” something about the value of the state θ . We refer to message m as DM1’s disclosure statement. This, as we discussed earlier, can take either oral or written form.

Obviously, if the disclosure statement fully revealed the state θ , then DM2 would have no reason to investigate or otherwise expend any effort to verify the information. However, the bias creates a misalignment in preferences over the second period policy. Hence, just as in the well known cheap talk environments, the message by DM1 will never be fully informative.

Knowing this in advance, the uninformed DM2 has a choice before any policies or disclosure statements are undertaken. She can choose to rely exclusively on DM1’s message and thereby forgo any attempt to verify the information. Alternatively, she can hold open the option to seek additional information later on after the policies and disclosure occurred.

Formally, we let $\pi \in \{0, 1\}$ denote the “commitment” choice of DM2. The choice $\pi = 1$ denotes the commitment by DM2 to forgo both a policy audit and state verification and hence to rely on m alone, while $\pi = 0$ denotes “keeping ones options open.” Choosing $\pi = 0$, DM2 chooses to exercise discretionary authority at a later date over what information she will seek.

If the uninformed agent opts for discretionary authority ($\pi = 0$), she subsequently has three choices, i.e., three methods of obtaining information after DM1’s decisions. DM2 can (i) still choose not to audit, and therefore rely only on the message (we denote this choice by NA for “no audit”), (ii) undertake a costless audit to verify the policy choice a_1 of DM1 (we denote this choice by PA, for “policy audit”), or (iii) undertake a costly audit to verify the state directly (we denote this choice by SV, for “state verification”). To simplify, we assume that both SV and PA are “perfect” in the sense that SV fully reveals the state θ , while PA fully reveals the first period policy choice a_1 . We let the choice across these three possibilities be denoted by d so that

$$d \in \{NA, PA, SV\} \tag{4}$$


Other things equal, DM2 would surely (weakly at least) prefer State Verification over a Policy Audit since the state, not the prior action, is relevant to her policy choice in period 2. On the other hand, the act of verifying the state directly may be difficult or costly in many common situations. Moreover, the first period policy, while having no direct effect on the

⁸ One example suggested to us by a referee concerns a common pool problem with an unknown resource stock. Initial users of the stock (e.g., fisherman) observe information about the stock (the fishing that day was good) which can be communicated or monitored (audited) by subsequent resource users. Agency bias arises when the “catch” is a purely private good, whereas ideological bias arises when the users have different preference intensities over a “shared catch.”

second period payoff, may nonetheless have signaling value for the decision makers. In other words, since θ affects DM1's preferences over his choice of a_1 , it is clear that a Policy Audit can reveal information about θ .

Of course, if SV is much more costly than PA, then it is not a viable option. To capture the idea that SV is only marginally more costly than SV, we assume that Policy Audits are costless while State Verification carries a lexicographic cost.⁹ The assumption of zero cost for PA is without loss of generality as it pertains to our main result. Specifically, the main result here refers the surprising outcome that DM2 may commit to disclosure and thus forego PA and SV. Adding a cost to PA therefore only reinforces this result. Moreover, if the PA cost is sufficiently small, all our results go through.¹⁰

The timing is depicted schematically below.



3.3. Equilibrium

We let $a_1 = \alpha_1(\pi, \theta)$ denote the choice of policy by DM1 given his observation of θ and DM2's choice of π . Similarly, we let $m = \mu(\pi, \theta)$ denote DM1's choice of message m . We also denote by $z(\pi, d)$ the actual observed outcome corresponding to the transmission choices π and d by DM2. According to the model, $z(1, d) = m$ for all d (i.e., DM2 chooses to commit to view only messages), $z(0, d) = m$ when $d = NA$, $z(0, d) = \theta$ when $d = SV$, and $z(0, d) = a_1$ when $d = PA$. Finally, let $a_2 = \alpha_2(z(\pi, d))$ denote DM2's choice of policy given her observation $z(\pi, d)$.

We evaluate key properties of the Perfect Bayesian Equilibria (PBE) of the game. For now we restrict attention to PBE in pure strategies, defined formally below. Aside from the simplification this affords, we show in Section 5 that mixed strategy PBE are non-generic in the bias parameter b . Moreover in those (non-generic) cases where mixed PBE may exist, one can construct a pure strategy PBE that is payoff equivalent for DM2.

In the definition below, we will refer to a PBE is simply an "equilibrium."

Definition 1. *Equilibrium:* An equilibrium is a 5-tuple $(\alpha_1, \mu, \pi, d, \alpha_2)$ such that

- (1) DM1 chooses the pair (α_1, μ) optimally given her observation of θ and π and DM2's (correctly anticipated) choice of d and α_2 .
- (2) DM2 chooses the triple (π, d, α_2) optimally. In particular, π and d are chosen given her beliefs about θ and her correct anticipation of DM1's choice of α_1 and μ . DM2's choice of α_2 is similarly optimal, reflecting her updated beliefs after observing $z(\pi, d)$.
- (3) DM2's updated beliefs satisfy Bayes' Rule wherever possible.

We first establish a baseline partial characterization of equilibria that will be helpful for subsequent results.

Proposition 1. *Full revelation without commitment:* Consider any equilibrium in either the ideological or the agency bias model. Suppose that in the equilibrium DM2 chooses $\pi = 0$. That is, DM2 does not commit to observing only message m and instead retains the option to conduct a policy audit or verify the state at a later stage. Then

- (i) DM1's choice of $a_1 = \alpha_1(0, \theta)$ is fully revealing of the state θ .
- (ii) DM1's message strategy $\mu(0, \cdot)$ in the continuation game following $\pi = 0$ coincides with some PBE strategy in the CS cheap talk game.
- (iii) DM2 chooses a policy audit so that $d = PA$ and $z(0, PA) = \alpha_1(0, \theta)$.
- (iv) DM2 chooses her optimal second period action $a_2 = \alpha_2(z(0, PA)) = \alpha_2(\alpha_1(0, \theta)) = \theta$.

One significant implication of Proposition 1 is that in a world where the commitment to forgo policy auditing and state verification is not possible, the uninformed agent would always choose to undertake a policy audit. She would never rely on messages alone, and she would never choose to verify the state directly.

In the interest of brevity, we omit a full formal proof for Proposition 1. The argument is standard and what follows is a sketch of the steps involved.

⁹ That is, PA is costless to implement by DM2, and DM2 chooses SV if and only if the payoff after observing the state is strictly larger than the payoff after observing DM1's action.

¹⁰ The argument is made more precise in Section 5.

Assume that in some equilibrium, $\pi = 0$. When DM2 chooses d it is not possible that she will set $d = NA$ and only observe m . In this case, given that a_1 is sunk, the model would boil down to a version of the CS world, and the message would not be fully revealing of the state. Hence DM2 could deviate to setting $d = SV$, observe θ , and hence attain a higher payoff.

Could it then be that DM2 sets $d = SV$ in equilibrium? In this case DM1 would anticipate that his choice of a_1 plays no (informational) role in DM2's choice of a_2 . Hence DM1 would simply choose an action that maximizes his $\tau = 1$ payoff, and hence his choice of a_1 would be fully revealing of the state. But since both SV and PA would induce fully revealing actions by DM1 and are thus equally informative, the policy audit PA is preferable to DM2 as the cheaper option. Hence, DM2 would then deviate to choosing PA.

This in turn means that only possibility is that DM2 chooses PA when given the chance. Could it be that this is the case and DM1's choice of a_1 is not fully revealing of θ ? If this were the case, given that lexicographic cost of choosing SV, DM2 would then gain by deviating to choosing SV. This would give her more information and hence allow her to improve her payoff. Hence, we conclude that in any equilibrium if $\pi = 0$ is chosen then $d = PA$ is also chosen by DM2, and furthermore that the choice of a_1 on DM1's part must be fully revealing of the state θ .

From Proposition 1, the only candidates for equilibrium are either for DM2 to either rely on the "words" (the case of $\pi = 1$) or the "actions" (the case of $\pi = 0$ and $d = PA$) of DM1. Notice that while SV will never be chosen in equilibrium, its availability is crucial to the argument. Without it, $d = NA$ could be chosen by DM2 and, consequently, the policy of DM1 need not be fully revealing.

4. When do actions speak louder than words?

4.1. Ideological bias

We begin with the case of ideological bias. In this case DM1 and DM2 differ in their objectives in the same way across the two periods. For any given θ DM1's ideal policy is larger than that of DM2. As we remarked before, this type of misalignment on preferences is a stylized representation of political competition between two parties that hold power in sequence but have a pre-set difference of opinion regarding the desirable values of a given policy variable.

From Proposition 1 we know that in the continuation of the game following $\pi = 0$ (no commitment), DM2 chooses $d = PA$ and DM1's choice of a_1 must be fully revealing of the state θ . Given that in cheap talk models with a continuum of states (à la CS) fully informative equilibria are generally ruled out, the picture that emerges from Proposition 1 seems to say that in equilibrium we should not observe DM2 actually commit to only observing the cheap talk message m .

Yet, this intuition turns out to be false in a robust set of environments. The main result tells us that with Ideological Bias and for "large enough" values of b , every equilibrium will have the property that DM2 commits to observe the messages. The gist of the argument relies on finding upper and lower bounds for DM2's payoff following a choice of $\pi = 0$ and $\pi = 1$ respectively. This is the reason we first state a preliminary characterization of equilibria, which sharpens the conclusion of Proposition 1 in the case of Ideological Bias.

Proposition 2. *Linear pareto-dominant equilibrium: Consider the Ideological Bias case. If there is any equilibrium in which DM2 chooses $\pi = 0$, then there is an equilibrium as follows.*

- (i) DM2 chooses $\pi = 0$.
- (ii) DM1's choice of policy is given by $\alpha_1(0, \theta) = \theta + 2b$.
- (iii) The payoffs associated with this equilibrium Pareto-dominate those of all other equilibria in which DM2 chooses $\pi = 0$.

For lack of a better term, we will refer to the equilibrium described in the Proposition as the linear equilibrium under Ideological Bias. We are now ready to state our first main result.

Proposition 3. *Sometimes words do speak louder than actions: Consider the Ideological Bias case. Then there exists a $\bar{b} > 0$ such that*

- (i) If $b > \bar{b}$, then in any equilibrium DM2 selects $\pi = 1$. In other words, in all equilibria DM2 commits to observing m only and forgoes the option to conduct a policy audit or to verify the state at a later stage.
- (ii) If $b \leq \bar{b}$, then there exists an equilibrium in which DM2 selects $\pi = 0$, subsequently chooses to audit DM1's policy (selects $d = PA$) and DM1 plays according to the linear equilibrium of Proposition 2, thus setting $a_1 = \alpha_1(0, \theta) = \theta + 2b$.

Obviously, the stronger and more surprising claim in Proposition 3 is part i which applies to every equilibrium of the game when $b > \bar{b}$. Thus, with a large enough ideological bias DM2 commits to observe only m . Words speak louder than actions.

We emphasize that the commitment option to forego a discretionary audit appears in many environments, including national security issues and standard corporate practices. The Proposition provides one rationale for why the option will be exercised. We later show that the result holds up even if the commitment by DM2 required the consent of DM1. For bias large enough, both parties would prefer to avoid the audit (see Section 5).

The full Proposition 3 is in Appendix. To understand the result, however, we find it useful to first develop the intuition for part i. Later, we comment on part ii.

Since we are in the case of ideological bias, for any given level of θ , DM1's ideal actions are $\theta + b$ both in period $\tau = 1$ and $\tau = 2$. The second decision maker DM2 favors actions equal to θ in both periods.

As the "receiver" of information, DM2 chooses to set $\pi = 1$ if b is large enough. In so doing, DM2 commits to observe message m only. Note that DM2 loses valuable information since, just as in CS, m is *not* fully informative of θ . In fact the degree of informativeness of m decreases as the size of b increases.

After the DM2 chooses $\pi = 1$, DM1 effectively chooses a_1 "in secret." Hence DM1 will simply set $a_1 = \theta + b$. Therefore, regardless of θ the first period payoff of DM2 is $-b^2$ (see (1)). DM2 then chooses a_2 after observing m , and hence (using (1) again) her second period expected payoff contingent on m is $-\text{Var}(\theta|m)$.¹¹

If DM2 were to set $\pi = 0$ instead, we know from Proposition 1 that it will subsequently choose to set $d = PA$, and hence observe a_1 directly. The key to seeing that, for large enough b , this will be worse than setting $\pi = 1$, is to note that after DM2 sets $\pi = 0$, DM1 knows that a_1 will be audited by DM2 in the second period (again we observe the fact that SV is available and dominates NA is critical to this argument, even though SV is never chosen in equilibrium). By Proposition 1, DM2 will have full information about θ once it observes a_1 . Therefore DM2 sets $a_2 = \theta$, and this gives DM2 a second period payoff of 0.

Knowing that DM2 will use the information yielded by the Policy Audit to set $a_2 = \theta$ will induce DM1 to "over-compensate" due to the marginal effect of its first period action on DM2's action. DM1 therefore sets a value of a_1 that is above $\theta + b$. In fact, we know from Proposition 2 that in the best case for DM2 (and, for that matter, for DM1 as well), a_1 will be set to equal $\theta + 2b$. Hence the first-period payoff to DM2 is now worse than the one we calculated following $\pi = 1$ since it is $-4b^2$.¹² Since, as we noted above, the second period payoff to DM2 is zero, this means that its overall expected payoff after setting $\pi = 0$ is at best $-4b^2$.

To sum up, DM2's expected payoff after setting $\pi = 1$ is bounded below by $-b^2 - \text{Var}(\theta)$, while after setting $\pi = 0$ it is bounded above by $-4b^2$. Since the distribution of θ is given, it is then clear that when b is sufficiently large DM2 must prefer to set $\pi = 1$ and thus forgo any Policy Audit or State Verification and rely only on the cheap talk message m sent by of DM1. Words *do* speak louder than actions in this case. In fact, it is not hard to extend the argument used to prove Proposition 3 to show that if b is sufficiently large, then DM2's commitment to observe only m is beneficial to *both* players.

The statement in Part (ii) of Proposition 3 is not as sharp as one would ideally like because of the underlying multiplicity of continuation equilibria following both $\pi = 0$ and $\pi = 1$ in the case where b is small.

The equilibrium identified in Part (ii) of Proposition 3 is supported (see Appendix for the details of the argument) by choosing a babbling equilibrium continuation following a choice of $\pi = 1$ and the linear (Pareto-dominant) equilibrium identified in Proposition 2 following a choice of $\pi = 0$.

In the model we analyze it is DM2 who chooses whether the game continues with an option to perform Policy Audit or State Verification or to shut down these options and commit to observing m only. It is therefore legitimate to ask what happens if, in the spirit of forward induction arguments,¹³ we restrict attention to "continuation equilibria" that give the best possible payoff to DM2 *both* after a choice of $\pi = 0$ and after a choice of $\pi = 1$ (the latter is not a babbling equilibrium).

Definition 2. *Forward induction proof equilibrium:* Consider an equilibrium of the general game. We define a *forward induction proof equilibrium* (henceforth a FIPE) as a PBE such that there is no other PBE which yields a larger continuation payoff to DM2 following either of its commitment actions, either $\pi = 0$ or $\pi = 1$.

In other words, a PBE is a FIPE if it maximizes, among all PBE's, DM2's continuation payoff following her commitment decision. In the spirit of forward induction, DM2 indicates her intentions about which action she'll take from her initial commitment decision. Off-path beliefs of DM1 accepts these intentions and acts accordingly.

Before stating the result, we require one more assumption.

Assumption 1. *Constant density:* The strictly positive continuous density f from which θ is drawn is constant on some non-degenerate interval $[\underline{\theta}, \bar{\theta}]$.

Proposition 4. *Consider the Ideological Bias case and let Assumption 1 hold. Then there exists a $\underline{b} > 0$ such that for any $b \leq \underline{b}$ the model has a unique FIPE in which DM2 sets $\pi = 0$, and subsequently chooses $d = PA$.*

Clearly, $\underline{b} < \bar{b}$ where \bar{b} is the bound in Proposition 3. So, with the refinement and a small enough ideological bias, a policy audit must occur. Combining this result with part (i) of the previous proposition, we obtain a sharp characterization of the chosen vehicle for information transmission when ideological bias is either sufficiently large or sufficiently small.

¹¹ The actual value of $\text{Var}(\theta|m)$ depends on the entire message function $\mu(\cdot|0)$ chosen by DM1 in equilibrium.

¹² Since we know from Proposition 1 that a_1 is fully revealing of the state θ , we can write the incentive-compatibility constraint for DM1 as balancing the marginal "tangible" payoff from a_1 (as it enters directly DM1's first period payoff) and its marginal effect on the beliefs of DM2 about θ and hence on a_2 . This gives rise to a simple differential equation (see (.5) in Appendix) which underpins the linear equilibrium of Proposition 2.

¹³ For reasons of space and ease of reading, we stay clear of details about how existing forward induction refinements specifically apply to our model. Instead, we refer the interested reader to the contributions by van Damme (1989), and more recently by Govindan and Wilson (2009), and Man (2012).

4.2. Agency bias

Our model behaves quite differently when the mis-alignment in preferences between DM1 and DM2 takes the form of Agency Bias. Recall (see (3)) that in this case DM1 favors a higher level of a_1 than DM2 does, and similarly, DM2 favors a higher level of a_2 than DM1 does. The manager who is active in any period, for any given θ (the external market factors affecting the firm), favors a higher level of managerial perks than in the period in which she is not the active manager but holds, say, stock in the firm.

Just as in the case of Ideological Bias, from Proposition 1 we know that only two alternatives are viable in equilibrium. Either DM2 chooses to *commit* to observing only m and thus forgoes any possibility to audit DM1's policy or to verify the state (DM2 sets $\pi = 1$), or she does not commit in such a way (DM2 sets $\pi = 0$) and subsequently chooses to observe a_1 ($d = PA$). Recall also that, just as in the case of Ideological Bias, Proposition 1 further characterizes this second possibility telling us that in this case DM1's choice of a_1 must be fully revealing of the state θ .

We first state our main result for the case of Agency Bias and then elaborate on its meaning and the intuition behind it.

Proposition 5. *Under agency bias actions speak louder than words: Consider the case of Agency Bias. Then, for any level of b , there is a unique Forward Induction-Proof equilibrium (FIPE) in which*

- (i) DM2 chooses $\pi = 0$ thereby not committing to observing m alone and retains the option to carry out a Policy Audit or State Verification.
- (ii) DM2 chooses a policy audit so that $d = PA$ and $z(0, PA) = \alpha_1(0, \theta)$.
- (iii) DM2 enjoys her globally optimal payoff of 0.

Note that combining Proposition 5 with Proposition 1 we know that in the equilibrium singled out by Proposition 5, DM2 proceeds to set $d = PA$ and so to audit DM1's policy and observe a_1 directly, thus gaining full knowledge of the actual value of θ .

Just as in the case of Proposition 3, there may be a multiplicity of PBE. However, FIPE rules out any continuation other than the one in which (a) the fully revealing (linear) equilibrium is played following the choice of $\pi = 0$, and (b) the best CS equilibrium (from DM2's point of view) is played following the choice of $\pi = 1$. Faced with these two continuations, DM2 will choose $\pi = 0$ in the agency bias model.

The intuition for the result can be understood by juxtaposing the agency bias model with the ideological bias model. In both models, the action of the first decision maker, DM1, plays two roles. Firstly, it has "tangible" value for both DM1 and DM2 since it enters directly their payoffs. Secondly, it has a signaling role since it potentially reveals the actual value of θ to DM2 – in fact, by Proposition 1, it actually *does* reveal it in any equilibrium in which $\pi = 0$. In both models, the marginal effect (via signaling) of DM1's action on DM2's action will be weighted by DM1 together with the tangible value of her action as it enters her first period payoff.

Unlike the ideological bias model, under agency bias DM1 would like to distort DM2's action downward rather than upward. In order to do this via signaling, his own action should be less distorted than would be optimal from a purely tangible payoff point of view. That is to say, the signaling value of DM1's action works in opposition to its tangible value. This is quite different than the ideological bias case where tangible value and signaling value are reinforcing – both motivate DM1 to distort his action upward.

Under agency bias the two opposing forces effectively cancel out making DM1's choice equal to the one that is the overall favorite of DM2, namely $a_1 = \theta$.¹⁴ Hence DM2 ends up enjoying her overall optimal payoff under a Policy Audit that is anticipated by DM1. Hence she chooses *not to commit* to observing m only when given the chance. Actions do indeed speak louder than words in the Agency Bias case.

5. Robustness

A number of extensions or "robustness checks" are worth discussing. We introduce what we believe are the most critical ones and analyze each in turn.

Mixed strategies. The analysis thus far focussed exclusively on pure strategies. We give a brief argument demonstrating the main results extend to the mixed strategy case, although we do not give a precise characterization of all mixed equilibria of the model.

Consider first Proposition 1, the precursor to subsequent results. Suppose we are at the stage in which DM2 did not already commit to disclosure and must now choose between PA (auditing), NA (messages), and SV (state verification) following DM1's choices.

Is there an equilibrium in which DM2 mixes at this stage? Notice at the outset that DM2 must assign zero probability to NA. If DM2 assigns positive probability to NA, then DM1's message is not fully informative (as shown in CS). In that case, DM2 strictly prefers SV to NA since the verification cost is smaller than its benefits. Notice also that DM2 will not randomize

¹⁴ The differential equation embodying DM1's Incentive Compatibility constraint (see (.17) in Appendix) now yields $a_1 = \theta$ as one of the solutions.

between PA and SV. If DM1's action is fully revealing, then DM2 strictly prefers PA. If the action is not fully revealing then DM2 strictly prefers SV. Hence, in equilibrium DM2 chooses PA.

Will DM2 actually be willing to randomize at the commitment stage? In fact, this can happen only if DM2 is indifferent between PA and NA. In the ideological bias case with b large, babbling is strictly better for him than the fully revealing continuation.¹⁵ Similarly, with either agency bias, or ideological bias with b small, the fully revealing continuation is strictly better than any other.

In short, indifference by DM2 is non-generic. Moreover, if there does exist an equilibrium in which DM2 mixes, there is a pure strategy equilibrium that is payoff equivalent.

Auditing costs. The model assumes that auditing costs are zero. Consider what happens when we allow for positive auditing costs in the model. The results hold up if auditing costs are introduced, provided they are small enough. The key observation is that, as long as state verification is marginally costlier than the policy audit, the latter will always induce fully revealing actions by DM1. In that case, the policy audit will always be the better option. Consequently, [Proposition 1](#), the gateway to the subsequent results, continues to hold. Moreover, if the PA cost is small enough, then it will still be the preferred option over disclosure in those cases where PA was chosen in the model with zero auditing costs.¹⁶

Notice that if mixed strategies are introduced in the model with positive auditing costs, they will be nongeneric for the same reasons outlined above. Note also that if $\pi = 1$ is chosen under zero PA costs, then it will be chosen when such costs are positive. The more subtle case is when $\pi = 0$ and $d = PA$. Since the fully revealing continuation under PA is strictly better than any other when auditing costs are zero, it will be so when costs are low enough to preserve the strict inequality.¹⁷

Discounting. In a dynamic model, it is natural to ask what happens if the players discount the future, so that in our case the second period payoff so that it is multiplied by a factor $\delta \in (0, 1)$. The gist of our results survives intact the discounting of second period payoffs by DM1 and DM2.

[Propositions 2 and 3](#) continue to hold with the only amendment that the linear equilibrium is now $\alpha_1(0, \theta) = \theta + (1 + \delta)b$. The thresholds \bar{b} and $\underline{b} < \bar{b}$ need to be recalculated for the discounting case, but their existence remains exactly as stated in [Propositions 3 and 4](#). [Proposition 5](#) part i also continues to hold.

In short the overall picture framed by [Propositions 3–5](#) continues to hold when discounting is allowed. Words speak louder than actions in the case of Ideological Bias and b large, while actions speak louder than words when b is small and we are in the Ideological Bias case, or for any value of b when we are in the case of Agency Bias.

The ability to commit. We remarked earlier that commitments to forego audits are common practice in firms. To a more limited extent, they also exist in politics – particularly in areas of national security. One could take issue with the fact that, unlike our model where DM2 makes the commitment, these commitments are institutional in nature. They are often agreed upon either by non-partisan committees (e.g., the firm's Board of Directors), or unanimous agreement by the parties themselves. Our results can be extended to accommodate this possibility.

Consider an extra step in the model whereby a proposed commitment to forego auditing by DM2 requires the approval of DM1 – a Pareto criterion of sorts that determines whether the commitment to use disclosure is taken. Now consider two institutions (formally, two distinct games). In one game, the one we've analyzed so far, there exists a decision node in which DM2 can commit to use disclosure. In the other game, no such commitment option is available. That is, without the commitment option, DM2 must simultaneously choose between all three options – disclosure, auditing, direct verification after DM 1's action.

In the ideology bias model with the commitment option, when bias b is large, [Proposition 3](#) shows that DM2 prefers to commit to disclosure. But this is also true for DM1. In other words, for b large *both* decision makers prefer to make the commitment. To see why, observe that under auditing, DM1's total payoff is $-2b^2$. Under disclosure, the worst can happen to DM1 is babbling and DM1's lower bound in this case is $-Var(\theta) - b^2$. Thus for b large, disclosure – the commitment to avoid auditing – Pareto dominates the option to audit. In this sense, the two agents prefer an institution in which a commitment option exists to avoid discretionary audits.

By contrast, when b is small, then both parties prefer auditing. In this case the presence or absence of the commitment option in the model does not affect the outcome. The parties are therefore indifferent about which of the two games, i.e., which institutional environment, they are in.

Now consider the agency bias model with the commitment option. Here, when the bias b is small, then both parties would prefer auditing. Again there is a clear Pareto ranking but the presence or absence of the commitment option in the model does not affect the outcome. The parties are once again indifferent about which of the two games, i.e., which institutional environment, they are in.

For the sake of completeness, we observe that when b is large in the agency bias model, DM2 and DM1 have opposing preferences on the commitment decision. By [Proposition 4](#), the decision in the hands of DM2 results in auditing.¹⁸

¹⁵ We emphasize that DM1's choice depends on the *outcome* of DM2's mixed strategy, not on the mixed strategy itself. For this reason, DM1 does not mix, and so neither does DM2 except on a nongeneric set of b .

¹⁶ Roughly, we can bound the payoff of the most informative cheap talk equilibrium of the CS cheap talk game. Since the action by DM1 in the PA continuation is fully revealing, it will be strictly preferred if the auditing cost is very small.

¹⁷ There may of course be an auditing cost to create an indifference between PA and NA, but this case would be non-generic.

¹⁸ DM1 prefers commitment, and so the outcome of a planner's criterion would depend on the relative weights assigned to the two agents' payoffs.

Hence, except for this last case, both players would prefer, at least weakly, to be in an institution (a game) in which the commitment option – the option to forego discretionary auditing – is available.

Quadratic loss payoffs. Throughout, we have worked with preferences represented by time-separable quadratic-loss functions. Since CS, quadratic preferences have been the canonical benchmark case throughout the literature on cheap-talk. Across this field, quadratic preferences also have the privileged position of guaranteeing tractability, yielding explicit solutions in many cases. They play essentially the same role in our model.

It is legitimate to ask how crucial their role is in the qualitative flavor of the results we obtain here. We do not demonstrate a general result in the paper, however, the nature of the results indicate that these preferences are not “knife-edged.”

A simple way to see this is to observe that DM2’s transmission choices are discrete and generically her best responses in equilibrium are strict. This means that the arguments should go through unchanged in a well defined continuous parametric neighborhood of the quadratic preferences. To take one example, suppose $W_i = \phi V_i + (1 - \phi)U_i$ where V_i is the quadratic payoff specified before and U_i is a smooth function of actions, state and bias and strictly concave in the action. For the result to extend, one would show that the results remain intact for W_i if ϕ is close enough to one.¹⁹ We strongly conjecture this is the case, but do not prove it formally.

6. Audits and partisan investigations

This paper has emphasized the importance of endogenous choice of transmission protocol in decentralized decisions problems. There are many such problems in which the interests of relevant decision makers do not coincide. We focus on two cleavages – ideological and agency – that differentiate authority figures at different decision dates.

The results suggest that the commitment to a system of disclosure is more valuable when the cleavage between decision makers is ideological. We argued that, broadly speaking, the Ideological Bias case fits better the political economy scenario in which two parties successively hold power, while the agency case is arguably better suited to represent private organizations in which two decision makers are chosen internally to head the operations in succession.²⁰

This, in our view, tallies significantly with the fact that Auditing does in fact prevail in the case of firms. Audits (internal and external) are almost always thought of as “good management practices” in the case of the firms,²¹ whereas audits in the political sphere are often viewed as “partisan witch hunts.” When a new political party occupies power it has wide powers to renege on previous promises, if necessary by changing the law. One notable exception is in the area of national security. Institutional commitments to avoid public audits of sensitive security exist and are usually agreed upon by all parties.

It is worth emphasizing that decision problems of this nature also arise in more amorphous settings.²² Whether the institution is formal or informal, we point out that if legal/institutional commitments are feasible, then both decision makers at least weakly prefer them in all cases, agency or ideological bias, large or small b , except under agency bias when b is large.

Finally, the results indicate that when such commitment is not possible the outcome will in fact be an Audit of the actions of the previous party by the new holder of power. At this point the loss of welfare relative to the disclosure regime may well become apparent to the participants, but it simply is too late to avoid it. A future exploration into when, how, and why such commitments are possible in some environments and not in others would be well worth the effort.

Appendix.

Proof of Proposition 2. Consider the Ideological Bias case, and assume that there is an equilibrium with $\pi = 0$. By Proposition 1 then there is an equilibrium with $\pi = 0$, $d = PA$ and $\alpha_1(0, \theta)$ fully revealing of θ . Let this equilibrium be denoted by $(\alpha_1^*, \mu^*, \pi^*, d^*, \alpha_2^*)$. For every value of θ , DM1’s policy choice $a_1 = \alpha_1^*(0, \theta)$ must solve

$$\max_{a_1} -(a_1 - \theta - b)^2 - [\alpha_2^*(a_1) - \theta - b]^2$$

Assuming α_2^* is differentiable,²³ we can write the first order condition as

$$\alpha_1^*(0, \theta) - \theta - b + [\alpha_2^*(\alpha_1^*(0, \theta)) - \theta - b] \alpha_2^{*'}(\alpha_1^*(0, \theta)) = 0$$

Since $\alpha_1^*(0, \theta)$ is fully revealing and DM2’s optimal a_2 is equal to θ for any given θ , we must have that $\alpha_2^*[\alpha_1^*(0, \theta)] = \alpha_1^{*-1}[\alpha_1^*(0, \theta)] = \theta$. From the Inverse Function Theorem, we then know that $\alpha_2^{*'}[\alpha_1^*(0, \theta)] = 1/\alpha_1^*(0, \theta)$.

¹⁹ The exact cutoffs in b would obviously be different.

²⁰ These representations are, of course, approximate. For instance, office-motivated politicians – as opposed to the ideology-driven ones we have in mind here – may fit better the case of Agency Bias.

²¹ See Arter (2000).

²² See footnote 8.

²³ For this step of the argument of course it is sufficient to verify ex-post that function we find is indeed differentiable. In the proof of part iii) below this is not a viable option. However, since we know from Proposition 1 that we are dealing with fully revealing equilibria it is easy to see that DM1’s decision must be continuous and strictly monotone in θ and hence differentiable almost everywhere. This is sufficient for our purposes below, though the details are omitted for the sake of brevity.

Hence, the first order condition implies that the following differential equation must be satisfied

$$[\alpha_1^*(0, \theta) - \theta - b]\alpha_1^{*'}(0, \theta) - b = 0 \tag{.5}$$

Notice that $\alpha_1^*(0, \theta) = \theta + 2b$ is one solution to this differential equation (and, point by point, satisfies the required second order condition).

This suffices to prove part ii of Proposition 2. For future reference, at this point we notice that $\alpha_1^*(0, \theta) = \theta + 2b$ and $a_2 = \theta$ for every θ imply an equilibrium payoff for DM2 of $-4b^2$.

To show parts i and iii of Proposition 2, using part i of Proposition 1, it is sufficient to show that the payoffs associated with the linear equilibrium in which $\pi = 0$ and $\alpha_1^*(0, \theta) = \theta + 2b$ Pareto-dominate those of any other fully revealing equilibrium in which DM2 sets $\pi = 0$.

Consider again the differential equation in (.5), which we know any fully revealing equilibrium to be considered here must satisfy (almost everywhere).

Define

$$g(\theta) = \alpha_1(0, \theta) - \theta - b \tag{.6}$$

and notice

$$g'(\theta) = \alpha_1^{*'}(0, \theta) - 1 \tag{.7}$$

Hence, we can rewrite (.5) as

$$g(\theta) [1 + g'(\theta)] = b \tag{.8}$$

Then $g^*(\theta) = b$ is one solution to (.8), giving rise to the linear equilibrium strategy $\alpha_1^*(0, \theta) = \theta + 2b$ that we identified above.

In general, however, there are infinitely many solutions to (.8), corresponding to infinitely many fully revealing equilibria.²⁴

Once we have a g that solves (.8), for any given θ the payoffs to the decision makers in the corresponding equilibrium, can be respectively written as

$$V_1 = -g(\theta)^2 - b^2 \quad \text{and} \quad V_2 = -[b + g(\theta)]^2$$

Moreover, observe that, whenever $g(\theta) > 0$, both payoffs are decreasing in $g(\theta)$.

The rest of the proof consists of showing that if g is a solution to the differential equation in (.8) and $g \neq g^*$, then $g(\theta) \geq b$ for every θ . This will suffice to show that in any other equilibrium to be considered the payoffs to both decision makers are lower than those on the linear equilibrium, as required.

Consider again the differential equation in (.8). Begin by noting that directly from the differential equation itself we know that it is impossible that $g(\theta) = 0$ for any θ . Hence, since g is continuous and differentiable almost everywhere, if $g(\hat{\theta})$ is positive (negative) for some $\hat{\theta}$ then $g(\theta)$ must be positive (negative) for all $\theta \in \mathbb{R}$.

Now, by way of contradiction, consider a solution \tilde{g} to (.8) such that for some $\hat{\theta}$ we have $\tilde{g}(\hat{\theta}) < b$. We need to distinguish between two cases – first $\tilde{g}(\hat{\theta}) > 0$, and then $\tilde{g}(\hat{\theta}) < 0$.

Consider the case of $\tilde{g}(\hat{\theta}) > 0$ first. Since by our contradiction hypothesis $\tilde{g}(\hat{\theta}) < b$, directly from (.8) we get that $\tilde{g}(\theta) < \tilde{g}(\theta') < b$ for all pairs $\{\theta, \theta'\}$ satisfying $\theta < \theta'$ and $\theta' < \hat{\theta}$.

Restricting attention to an appropriate subsequence if necessary, we then have that $\lim_{\theta \rightarrow -\infty} \tilde{g}(\theta) = 0$. Since $\tilde{g}(\theta) < b$ for all $\theta < \hat{\theta}$, this implies that for some sufficiently small $\bar{\theta}$ it must be that

$$\tilde{g}(\bar{\theta})[1 + \tilde{g}'(\bar{\theta})] < b \tag{.9}$$

which directly implies that \tilde{g} cannot be a solution to (.8).

We now turn to the second case – so, assume $\tilde{g}(\hat{\theta}) < 0$. Similarly to what we argued in the previous case, this implies that $\tilde{g}(\theta') < \tilde{g}(\theta) < 0$ for all pairs $\{\theta, \theta'\}$ satisfying $\theta < \theta'$ and $\theta' < \hat{\theta}$.

Just as in the previous case, and again restricting attention to an appropriate subsequence if necessary, we then have that $\lim_{\theta \rightarrow -\infty} \tilde{g}(\theta) = 0$. Since $\tilde{g}(\theta) < 0$ for all $\theta < \hat{\theta}$, this implies that for some sufficiently small $\bar{\theta}$ it must be that

$$\tilde{g}(\bar{\theta})[1 + \tilde{g}'(\bar{\theta})] \leq 0 \tag{.10}$$

which directly implies that \tilde{g} cannot be a solution to (.8). \square

²⁴ Another solution is, for example, $\hat{g}(\theta) = b[W(e^{1-\theta/b}) + 1]$ where W is the so-called Lambert W function which is implicitly defined by $\theta = W(\theta) e^{W(\theta)}$. The function \hat{g} satisfies $\hat{g}(0) = 2b$ (which can be interpreted as the initial condition). The function \hat{g} therefore gives rise to another fully revealing equilibrium in which the informed DM's policy is $\hat{\alpha}_1(\theta, 0) = \hat{g}(\theta) + \theta + b$. It is also easy to check that there are infinitely many other solutions g to the differential equation in (.8). Suppose \tilde{g} is a solution to (.8). Fix Δ and define \hat{g} by letting $\hat{g}(\theta) = \tilde{g}(\theta + \Delta)$ for every θ . Then \hat{g} is also a solution.

Lemma A.1. Consider the case of Ideological Bias. Suppose that DM2 sets $\pi = 1$ and DM2 chooses α_1 and μ optimally. Then by choosing $a_2(m) = E(\theta)$ for every m , DM2 achieves an overall expected payoff of $-b^2 - \text{Var}(\theta)$.²⁵

Proof. Given that DM2 chooses $\pi = 1$, DM1's choice of a_1 will not be observed by DM2. Hence, from (1) it is evident that it is optimal for DM1 to set $a_1 = \theta + b$ for every $\theta \in \mathbb{R}$. Hence DM2's first period executed payoff is $-b^2$. Again directly from (1), ignoring m and setting $a_2 = E(\theta)$, it is clear that DM2 can achieve a second-period expected payoff of $\text{Var}(\theta)$. \square

Lemma A.2. Consider the case of Ideological Bias. In any equilibrium in which DM2 sets $\pi = 0$ her overall expected payoff is bounded above by $-4b^2$.

Proof. From Proposition 1 we know that such equilibrium must entail a choice of $d = PA$, and must be fully revealing. Hence DM2 will be able to infer θ from observing a_1 . Therefore, from (1), DM2 will choose $a_2 = \theta$ for every θ , and hence achieve a second period expected payoff of 0. This, of course, is DM2's globally optimal expected payoff in the second period.

From Proposition 2 we know that DM2's overall payoff cannot be above the one she obtains in the linear equilibrium in which DM1 sets $a_1 = \theta + 2b$ for every $\theta \in \mathbb{R}$. This gives DM2 a first period expected payoff of $-4b^2$. \square

Proof of Proposition 3. Let $\bar{b} = \sqrt{\text{Var}(\theta)/3}$. Note that this trivially implies that for every $b > \bar{b}$ we must have that $-4b^2 < -b^2 - \text{Var}(\theta)$, and for every $b \leq \bar{b}$ we have that $-4b^2 \geq -b^2 - \text{Var}(\theta)$.

We prove part i first. Suppose, by way of contradiction, that $b > \bar{b}$ and that there exists an equilibrium in which DM2 chooses $\pi = 0$. By Lemma A.2, in such a putative equilibrium DM2's payoff is bounded above by $-4b^2$. By Lemma A.1, by deviating and setting $\pi = 1$, DM2 will achieve a payoff that is bounded below by $-b^2 - \text{Var}(\theta)$. This is evidently a profitable deviation from the putative equilibrium and hence the proof of part i is now complete.

Now for the proof of part ii. Let any $b \leq \bar{b}$ be given, and construct an equilibrium as follows. DM2 chooses $\pi = 0$, and following this choice, on-path, the linear equilibrium of Proposition 2 is played. If DM2 deviates and sets $\pi = 1$, a babbling equilibrium is played in which DM1 chooses a fixed \bar{m} regardless of θ and DM2 chooses $a_1 = E(\theta)$. DM2's payoff in the proposed equilibrium is $-4b^2$. If instead he deviates and sets $\pi = 0$ his payoff is $-b^2 - \text{Var}(\theta)$. Since $b \leq \bar{b}$ this is not a profitable deviation and hence the proof of part ii is now complete. \square

Preliminary results for the proof of Proposition 4. Before proceeding with the proof of Proposition 4, we let φ denote the probability that $\theta \in [\underline{\theta}, \bar{\theta}]$ as required by Definition 1.

Definition A.1. Bounded partitional: An continuation equilibrium of the game following $\pi = 1$ is bounded partitional if and only if DM1's strategy partitions the set of possible θ (namely \mathbb{R}) into disjoint intervals (over which the message does not change), and the length of such intervals is bounded away from zero.

Lemma A.3. For any $b > 0$, all continuation equilibria of the game following $\pi = 0$ are bounded partitional.

Proof. This is a straightforward adaptation of arguments in CS (to the case of $\theta \in \mathbb{R}$ rather than a bounded interval), and the details are omitted. \square

Lemma A.4. Suppose that Assumption 1 holds. Let a decreasing sequence of strictly positive bias parameters $\{b_k\}_{k=0}^\infty$ with $\lim_{k \rightarrow \infty} b_k = 0$ be given.

Consider any sequence of associated continuation equilibria of the game following $\pi = 1$, and let V_{2k}^* be the sequence of associated payoffs for DM2.

Assume that for every k the equilibrium strategy of DM1 partitions $[\underline{\theta}, \bar{\theta}]$ into at most two sets (at most two distinct messages are sent in equilibrium as θ ranges over $[\underline{\theta}, \bar{\theta}]$).

Then, there exists a \bar{k} such that for every $k \geq \bar{k}$, we have that $V_{2k}^* < -4b_k^2$.

Proof. Let a b be given. Consider a strategy for DM1 that entails sending either one or two messages as θ ranges over $[\underline{\theta}, \bar{\theta}]$. Suppose only one message is sent, say \hat{m} . Conditional on $\theta \in [\underline{\theta}, \bar{\theta}]$, DM2's payoff is then bounded above by $-\text{Var}(\theta | \theta \in [\underline{\theta}, \bar{\theta}])$. Suppose next that only two message are sent, say \hat{m} and $\hat{\hat{m}}$. Then the payoff to DM2 is bounded above by $-\text{Var}(\theta | \theta \in [\underline{\theta}, \bar{\theta}] \text{ and } m = \hat{m}) - \text{Var}(\theta | \theta \in [\underline{\theta}, \bar{\theta}] \text{ and } m = \hat{\hat{m}})$. Note that these bounds do not depend on b , but only on the distribution of θ . Let the best possible bound be denoted by $-\ell$ (with $\ell > 0$). Then if DM1 only sends one or two messages as θ ranges over $[\underline{\theta}, \bar{\theta}]$ the second period payoff to DM2 is bounded above by $-\ell \varphi$. This is because DM2's payoff conditional on $\theta \notin [\underline{\theta}, \bar{\theta}]$ is at most 0. Since in any continuation equilibrium of the game following $\pi = 1$, DM1 obviously chooses $a_1 = \theta + b$, DM2's payoff in any such equilibrium is $-b^2$. Hence, we can conclude that if DM1 only sends one or two messages as θ ranges over $[\underline{\theta}, \bar{\theta}]$ in any continuation equilibrium of the game following $\pi = 1$, the overall payoff to DM2 is bounded above by $-b^2 - \ell \varphi$.

Since $-\ell \varphi$ does not depend on b , for b sufficiently small it is evident that $-b^2 - \ell \varphi < -4b^2$. The claim then follows immediately. \square

Proof of Proposition 4. By way of contradiction assume that there exists a decreasing sequence of bias parameters $\{b_k\}_{k=0}^\infty$ with $\lim_{k \rightarrow \infty} b_k = 0$, such that for each b_k in the sequence the model has a FIPE in which DM2 actually chooses $\pi = 1$.

²⁵ Here $E(\theta)$ denotes the expected value of θ .

Notice first that by [Propositions 1 and 2](#) we know that in any FIPE the players' strategies conditional on $\pi = 0$ give rise to the linear equilibrium of [Proposition 2](#) and hence yield a payoff of $V_{2k}^* = -4b_k^2$ to DM2.

By [Lemma A.4](#), it must be that, contingent on $\pi = 1$, for every k in the sequence (discard the initial elements if necessary) DM1's strategy involves sending at least three different messages as θ ranges over $[\underline{\theta}, \bar{\theta}]$.

Our next endeavor is to compute an upper bound to DM2's payoff contingent on $\pi = 0$ – or equivalently in the continuation equilibrium of the game following $\pi = 1$.

For any given b_k , consider the partition of the set of possible θ (namely \mathbb{R}) induced by DM1's strategy (see [Lemma A.3](#)). Consider the left-most cell of such partition such that its upper bound is no smaller than $\underline{\theta}$. Let such bound be denoted by \underline{d}_k . Consider next the right-most cell of such partition such that its lower bound is no larger than $\bar{\theta}$. Let such bound be denoted by \bar{d}_k . Without loss of generality (by taking subsequences if necessary) we assume that both sequences $\{\underline{d}_k\}_{k=0}^\infty$ and $\{\bar{d}_k\}_{k=0}^\infty$ converge as $k \rightarrow \infty$.

Notice next that it must be the case that

$$\lim_{k \rightarrow \infty} (\bar{d}_k - \underline{d}_k) = \eta > 0 \tag{.11}$$

This is simply because otherwise we could obtain a contradiction using the same argument (in the limit) as in the proof of [Lemma A.4](#). Let ϕ_k be the probability that $\theta \in [\underline{d}_k, \bar{d}_k]$ and $\lim_{k \rightarrow \infty} \phi_k = \phi$. Note that by [Assumption 1](#) it must be that $\phi = \eta\varphi/(\bar{\theta} - \theta)$ and hence $\phi > 0$.

Notice that at this point we know that contingent on $\theta \in [\underline{d}_k, \bar{d}_k]$, the continuation equilibrium of the game following $\pi = 1$ is exactly the model in CS in the special case of a uniform density.²⁶

Let $U(\underline{d}_k, \bar{d}_k, b_k)$ be DM2's second period payoff under the most informative (the one that gives DM2 the highest payoff), contingent on $\theta \in [\underline{d}_k, \bar{d}_k]$. Overall, DM2's equilibrium payoff in the continuation equilibrium of the game following $\pi = 1$ is then bounded above by

$$-b_k^2 + \phi_k U(\underline{d}_k, \bar{d}_k; b_k) \tag{.12}$$

This is simply because DM2's first period payoff in any equilibrium of the continuation equilibrium of the game following $\pi = 1$ is $-b_k^2$ and whenever $\theta \notin [\underline{d}_k, \bar{d}_k]$ DM2's second period payoff obviously cannot exceed 0.

Our contradiction hypothesis and [\(.12\)](#) imply that there exists $n \in \mathbb{N}$ such that for every $k \geq n$,

$$-b_k^2 + \frac{\phi}{2} U(\underline{d}_k, \bar{d}_k, b_k) \geq -4b_k^2 \tag{.13}$$

which implies

$$U(\underline{d}_k, \bar{d}_k; b_k) \geq -\frac{6b_k^2}{\phi} \tag{.14}$$

Directly from CS in the constant density case²⁷ we get that

$$U(\underline{d}_k, \bar{d}_k, b_k) = -A_k b_k^2 - B_k \tag{.15}$$

for some positive A_k and B_k . Also directly from CS, there exists $m \in \mathbb{N}$ such that for every $k \geq m$ we have that $A_k > 6/\phi$. Hence, for $k \geq \max\{n, m\}$ it must be that

$$U(\underline{d}_k, \bar{d}_k, b_k) = -A_k b_k^2 - B_k < -\frac{6b_k^2}{\phi} \tag{.16}$$

which contradicts inequality [\(.14\)](#) and hence concludes the proof \square

We continue with a preliminary result that will be used to prove [Proposition 5](#).

Lemma A.5. Consider the Agency Bias case. If there is any equilibrium in which DM2 chooses $\pi = 0$, then there is an equilibrium as follows.

- (i) DM2 chooses $\pi = 0$.
- (ii) DM1's choice of policy is given by $\alpha_1(0, \theta) = \theta$.
- (iii) The payoffs associated with this equilibrium Pareto-dominate those of all other equilibria in which DM2 chooses $\pi = 0$.

We will refer to this as the linear equilibrium under Agency Bias.

²⁶ CS consider explicitly the case of a θ uniformly distributed over $[0, 1]$. An appropriate scale factor needs to be factored in to conform to the (sub-)model we have here.

²⁷ See footnote 26 above.

Proof. Consider²⁸ the Agency Bias case, and assume that there is an equilibrium with $\pi = 0$. By Proposition 1 then there is an equilibrium with $\pi = 0$, $d = PA$ and $\alpha_1(0, \theta)$ fully revealing of θ . Let this equilibrium be denoted by $(\alpha_1^*, \mu^*, \pi^*, d^*, \alpha_2^*)$. For every value of θ , DM1's policy choice $a_1 = \alpha_1^*(0, \theta)$ must solve

$$\max_{a_1} -(a_1 - \theta - b)^2 - [\alpha_2^*(a_1) - \theta]^2$$

Assuming α_2^* is differentiable,²⁹ we can write the first order condition as

$$\alpha_1^*(0, \theta) - \theta - b + [\alpha_2^*(\alpha_1^*(0, \theta)) - \theta]\alpha_2^{\prime}(\alpha_1^*(0, \theta)) = 0$$

Since $\alpha_1^*(0, \theta)$ is fully revealing and DM2's optimal a_2 is equal to $\theta + b$ for any given θ , we must have that $\alpha_2^*[\alpha_1^*(0, \theta)] = \alpha_1^*^{-1}[\alpha_1^*(0, \theta)] + b = \theta + b$. From the Inverse Function Theorem, we then know that $\alpha_2^{\prime}[\alpha_1^*(0, \theta)] = 1/\alpha_1^{\prime}(0, \theta)$.

Hence, the first order condition implies that the following differential equation must be satisfied

$$[\alpha_1^*(0, \theta) - \theta - b]\alpha_1^{\prime}(0, \theta) + b = 0 \tag{.17}$$

$\alpha_1^*(0, \theta) = \theta$ is one solution to this differential equation (and, point by point, satisfies the required second order condition).

This suffices to prove parts i and ii of Lemma A.5.

To show part iii of Lemma A.5, using ii of Proposition 1, it is sufficient to show that the payoffs associated with the linear equilibrium in which $\pi = 0$ and $\alpha_1^*(0, \theta) = \theta$ Pareto-dominate those of any other fully revealing equilibrium in which DM2 sets $\pi = 0$.

Consider again the differential equation in (.17), which we know any fully revealing equilibrium to be considered here must satisfy (almost everywhere), and define

$$g(\theta) = \alpha_1(0, \theta) - \theta - b$$

and notice

$$g'(\theta) = \alpha_1^{\prime}(0, \theta) - 1 \tag{.18}$$

Hence, we can rewrite (.17) as

$$g(\theta) [1 + g'(\theta)] = -b \tag{.19}$$

$g^*(\theta) = -b$ is one solution to (.19), giving rise to the linear equilibrium strategy $\alpha_1^*(0, \theta) = \theta$ that we identified above.

In general, however, there are infinitely many solutions to (.19), corresponding to infinitely many fully revealing equilibria.³⁰

Once we have a g that solves (.19), for any given θ the payoffs to the decision makers in the corresponding equilibrium, can be respectively written as

$$V_1 = -g(\theta)^2 - b^2 \quad \text{and} \quad V_2 = -[b + g(\theta)]^2$$

Moreover, observe that, whenever $g(\theta) < -b$, both payoffs are increasing in $g(\theta)$.

The rest of the proof consists of showing that if g is a solution to the differential equation in (.19) and $g \neq g^*$, then $g(\theta) \leq -b$ for every θ . This will suffice to prove that in any other equilibrium to be considered the payoffs to both decision makers are lower than those on the linear equilibrium, as required.

Consider again the differential equation in (.19). Begin by noting that directly from the differential equation itself we know that it is impossible that $g(\theta) = 0$ for any θ . Hence, since g is continuous and differentiable almost everywhere, if $g(\hat{\theta})$ is positive (negative) for some $\hat{\theta}$ then it must be that $g(\theta)$ is positive (negative) for all $\theta \in \mathbb{R}$.

Now, by way of contradiction, consider a solution \tilde{g} to (.19) such that for some $\hat{\theta}$ we have $\tilde{g}(\hat{\theta}) > -b$. We need to distinguish between two cases – first $\tilde{g}(\hat{\theta}) < 0$, and then $\tilde{g}(\hat{\theta}) > 0$.

Consider the case of $\tilde{g}(\hat{\theta}) < 0$ first. Since by our contradiction hypothesis $\tilde{g}(\hat{\theta}) > -b$, directly from (.19) we get that $-b < \tilde{g}(\theta) < \tilde{g}(\theta')$ for all pairs $\{\theta, \theta'\}$ satisfying $\hat{\theta} < \theta$ and $\theta < \theta'$.

Restricting attention to an appropriate subsequence if necessary, we then have that $\lim_{\theta \rightarrow \infty} \tilde{g}'(\theta) = 0$. Since $\tilde{g}(\theta) > -b$ for all $\theta > \hat{\theta}$, this implies that for some sufficiently large $\bar{\theta}$ it must be that

$$\tilde{g}(\bar{\theta})[1 + \tilde{g}'(\bar{\theta})] > -b \tag{.20}$$

which directly implies that \tilde{g} cannot be a solution to (.19).

²⁸ The argument here, mutatis mutandis, runs along lines similar to the proof of Proposition 2, hence the exposition of some of the steps here will be streamlined.

²⁹ See footnote 23 above.

³⁰ See footnote 24 above.

We now turn to the second case – so, assume $\tilde{g}(\hat{\theta}) > 0$. Similarly to what we argued in the previous case, this implies that $0 < \tilde{g}(\theta') < \tilde{g}(\theta)$ for all pairs $\{\theta, \theta'\}$ satisfying $\hat{\theta} < \theta$ and $\theta < \theta'$.

Just as in the previous case, restricting attention again to an appropriate subsequence if necessary, we then have that $\lim_{\theta \rightarrow \infty} \tilde{g}'(\theta) = 0$. Since $\tilde{g}(\theta) > 0$ for all $\theta > \hat{\theta}$, this implies that for some sufficiently large $\hat{\theta}$ it must be that

$$\tilde{g}(\hat{\theta})[1 + \tilde{g}'(\hat{\theta})] \geq 0 \quad (.21)$$

which directly implies that \tilde{g} cannot be a solution to (.19). \square

Proof of Proposition 5. We simply show that DM2 setting $\pi = 0$ followed by the linear equilibrium identified in Lemma A.5 constitutes a unique FIPE of the model.

The claim is in fact rather obvious. Since in this equilibrium $\alpha_1^*(0, \theta) = \theta$, it is clear that DM2 achieves her globally optimal payoff of 0. It is easy to show that DM2's payoff is bounded away from zero if he commits to observe DM1's message. \square

References

- Anderlini, L., Gerardi, D., Lagunoff, R., 2012. Communication and learning. *Rev. Econ. Stud.* 79, 415–450.
- Arter, D., 2000. Management auditing. *Qual. Dig.*, 73.
- Austen-Smith, D., Banks, J., 2000. Cheap talk and burned money. *J. Econ. Theory* 91, 1–16.
- Boin, A., McConnell, A., Hart, P., 2008. *Governing After Crisis: The Politics of Investigation, Accountability and Learning*. Cambridge University Press, Cambridge.
- Bolton, P., Dewatripont, M., 2005. *Contract Theory*. MIT Press, Cambridge, MA.
- Carrillo, J., Mariotti, T., 2000. Strategic ignorance as a self-disciplining device. *Rev. Econ. Stud.* 67, 529–544.
- Crawford, V.P., Sobel, J., 1982. Strategic information transmission. *Econometrica* 50, 1351–1431.
- Daughety, A., Reinganum, J., 2008. Communicating quality: a unified model of disclosure and signaling. *RAND J. Econ.* 73, 973–989.
- Daughety, A., Reinganum, J., 2010. Public goods, social pressure, and the choice between privacy and publicity. *Am. Econ. J.: Microecon.* 2, 191–221.
- Farrell, J., Rabin, M., 1996. Cheap talk. *J. Econ. Perspect.* 10, 103–118.
- Goldberg, A., 2010. Been there, done that: Partisan investigations solve little. *Politico*.
- Govindan, S., Wilson, R., 2009. On forward induction. *Econometrica* 77, 1–78.
- Institute for Internal Auditors, 2012. International standards for the professional practice of internal auditing. Research report.
- Karamychev, A., Visser, D., 2011. An optimal signaling equilibrium. Technical report, Erasmus University and the Tinbergen Institute.
- Kartik, N., 2007. A note on cheap talk and burned money. *J. Econ. Theory* 136, 749–758.
- Kartik, N., 2009. Strategic communication with lying costs. *Rev. Econ. Stud.* 76, 1359–1395.
- Krishna, V., Morgan, J., 2008. Cheap talk. In: Blume, L., Durlauf, S. (Eds.), *The New Palgrave Dictionary of Economics*, 2nd ed. McMillan, New York, NY.
- Man, P.Y., 2012. Forward induction equilibrium. *Games Econ. Behav.* 75, 265–276.
- Prat, A., 2005. The wrong kind of transparency. *Am. Econ. Rev.* 95, 862–877.
- Sobel, J., 2007. *Signaling Games*. UCSD, mimeo.
- Spector, D., 2000. Rational debate and one-dimensional conflict. *Q. J. Econ.* 115, 181–200.
- Spence, A.M., 1973. Job market signaling. *Q. J. Econ.* 90, 225–243.
- Townsend, R., 1979. Optimal contracts and competitive markets with costly state verification. *J. Econ. Theory* 21, 265–293.
- van Damme, E., 1989. Stable equilibria and forward induction. *J. Econ. Theory* 48, 476–496.