ISSN 2279-9362

Collegio Carlo Alberto

Do Actions Speak Louder than Words?

Luca Anderlini
Dino Gerardi
Roger Lagunoff

No. 355
June 2014

Carlo Alberto Notebooks

www.carloalberto.org/research/working-papers

Do Actions Speak Louder Than Words?*

Auditing, Disclosure, and Verification in Organizations

Luca Anderlini
(Georgetown University)

Dino Gerardi
(Collegio Carlo Alberto)

Roger Lagunoff
(Georgetown University)

June 2014

Abstract. We study the relative performance of disclosure and auditing in organizations. We consider the information transmission problem between two decision makers who take actions at dates 1 and 2 respectively. The first decision maker has private information about a state of nature that is relevant for both decisions, and sends a cheap-talk 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 is preferences reflects a bias in a decision maker's *own action* relative to that of the other — we call this an agency bias — then, in equilibrium, the second decision maker chooses to audit. Actions speak louder than words in this case.

When one decision maker prefers *all actions* to be biased relative to the other decision maker — we call this an ideological bias — then, if the misalignment is large enough, in equilibrium the second decision maker chooses disclosure. In this case words speak louder than actions.

While firms are usually characterized by agency bias, ideological bias is more common in political systems. Our results indicate that the ability to commit not to audit has value in the latter case. However such commitment is rarely feasible in the political sphere.

JEL Classification: C73, D63, D72, D74, H11.
Keywords: Auditing, Disclosure, Agency Bias, Ideological Bias.
Address for correspondence: Luca Anderlini, Georgetown University, 37^{th} and O Streets NW, Washington DC 20057, USA. luca@anderlini.net

*We wish to thank seminar participants at Johns Hopkins, North Carolina, PUC (Chile), 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.

1. Introduction

In any ongoing organization there are two primary ways in which information is revealed. One is *disclosure* - transmission of information by the informed parties to uninformed receivers. Disclosure may be obligatory, but its content is determined by the informed sender and may therefore be subject to manipulation. The other is *auditing* — the attempt by the uninformed party to investigate or verify the information in some fashion. In this paper disclosure takes the extreme form of cheap talk and auditing takes the extreme form of being able to observe exactly actions taken by others.

This paper studies the relative performance of disclosure and auditing in a simple set-up.

We consider a two-period policy setting framework in which the policy at each date is chosen by a distinct decision maker. 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*, and she can do so *before* DM1 takes any action or sends his message. Throughout we refer to this as the case of *disclosure*. The decision by DM2 of whether to make such commitment or not is then announced to DM1. If she chooses not to commit, DM2 retains the option to observe DM1's action — throughout, we refer to this as *auditing* — or even to observe the state (at a cost) — throughout, we refer to this last possibility as *verification*. Once period 2 arrives, DM2 "reaches power" and her action is then taken.

Despite its simplicity, we believe the model encapsulates a variety of scenarios in a stylized way. It is standard practice in a firm, for instance, 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 outside shareholders and subsequent generations of managers.

Governments operate in much the same way. 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.

In both firms and governments, misaligned preferences between temporally separated agents (e.g., two managerial teams or two political parties) distort incentives away from full information disclosure. 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.

Our main concern is to study under what conditions disclosure will be used, and when instead will we observe reliance on auditing, or even verification, in equilibrium. The answer is nuanced and depends on the particular characteristics of the organization. Somewhat surprisingly, we find that there are plausible conditions under which, in equilibrium, the uninformed party will choose to commit to rely solely on disclosure and forego all other options, even when auditing costs are zero and even when the state can in principle be observed at an arbitrarily small cost. Sometimes it is indeed the case that, counterintuitively, "words speak louder than actions," in fact even louder than direct verification of the information that drives the actions in the first place.

It is self-evident that most auditing and investigative practices attempt to verify the *policies* of a decision maker rather than the underlying state directly. This is often because direct state verification is costly or impossible. Consider for instance, the United States Securities and Exchange Commission (SEC) regulators' attempts to investigate insider trading. In most cases, the regulators cannot possibly uncover the underlying knowledge of traders. Instead, the SEC looks for unusual trading patterns and/or price movements prior to a firm's disclosure (Heim, 2011). Similarly, political audits of potential wrongdoing involve verification of a politician's behavior in order to assess "what did [he] know and when did he know it?"[1]

In the model, verification by DM2 of the hidden state is costly, but not impossible. All our results hold when this cost is vanishingly small (lexicographic), and this is what we actually assume throughout. We believe that this strengthens our results when we find that words speak louder than actions since it supplies a superior alternative, even though it is rarely feasible in practice.

Our first result shows that even when the cost of direct state verification is vanishingly small, DM2 will never choose this 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, then uninformed DM2 would

[1] The quote refers to the famous question asked of President Nixon by Senator Howard Baker during the Watergate Hearings, 1973, p. 1476.

switch to the policy audit that has no cost.[2] Consequently, effectively DM2 chooses between the "words" (disclosure) and the "actions" (policy audit) of her predecessor, DM1.

Ultimately, the choice between words and actions depends on the wedge in preferences between the organizations' members. We refer to this misalignment as the "bias" (following CS) and examine two canonical cases. In the case of *agency bias*, the active decision maker prefers a higher (or lower) policy than his counterpart, the passive stakeholder. Arguably, many private firms conform to this case. This form of bias contrasts with the case of *ideological bias* whereby the first period decision maker always prefers systematically higher (or lower) policies than his second period counterpart, regardless of which party chooses or which period the choice is made. The ideological biases arise naturally in politics.

With agency bias, DM2 will generally choose a policy audit. However, under large enough ideological bias, DM2 will choose to commit in advance to use *only* the disclosed information (DM1's cheap talk message) and to forgo all subsequent audit and state verification possibilities.

To gain some intuition for these results, we observe that equilibria without auditing will resemble those in standard cheap talk models (e.g. CS). That is, only partial revelation of information occurs in equilibrium, and this is bad for DM2.

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 auditing in favor of disclosure is not just preferable for the recipient, but it can be welfare enhancing for all parties.

Our results suggest a paradox. In firms (where agency bias is prevalent), commitment to forgo auditing may be possible though unnecessary because audits are always welfare enhancing. In political systems, audits often assume the form of a political investigation. In this case the commitment to forgo auditing is difficult or impossible, though it may be highly

[2]Obviously this argument relies on state verification being more expensive than auditing the policy, not on the audit being free and the state verification having a vanishing small cost.

desirable. Without such commitment, however, auditing will occur, and this may make all parties worse off.

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. Section 6 concludes with anbrief discussion of the accountability systems found in firms and polities. 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 the Appendix.

2. Related Literature

We examine the equilibrium choice of disclosure or auditing by the uninformed agent.[3] With some exceptions (discussed below) most of the literature examines either one or the other transmission mechanism, but not both. Disclosure is the main focus of the cheap talk communication literature beginning with CS.[4]

The present paper contains elements of dynastic cheap talk reminiscent of Spector (2000) and Anderlini, Gerardi, and Lagunoff (2012). Those papers examine the effect 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.

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 deeds" trade off somewhat resembles their "privacy versus public information" trade-off even though there is no (cheap talk) communication in the

[3]A possible course of action open to DM2 is verification, but, as we noted above, it is relatively straightforward to rule out this possibility in equilibrium.

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

latter.

A number of papers have looked at both cheap talk and signaling. 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), 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 f — a strictly positive continuous density over $\mathbb{R}$ — before anything else takes place. The four quantities $b_{t,\tau} \in \mathbb{R}$ with $t = 1, 2$ and $\tau = 1, 2$ are a set of "bias" parameters that differentiates the payoffs of the two DMs — we refer to $b_{t,\tau}$ as the bias of DMt in period τ.

Each DMt with $t = 1, 2$ has two period payoff

$$V_t \ = \ -(a_1 - \theta - b_{t,1})^2 \ - \ (a_2 - \theta - b_{t,2})^2 \tag{1}$$

Note that (1) implies that both DMs care equally about both periods in the chronology of the organization. In Section 5 we return to this point and note that our main results go through for arbitrary discount factor $\delta \in (0, 1)$ applied to second period payoffs.

To make our main points it suffices to consider only two possible values for each of the four parameters $b_{t,\tau}$. Each of the $b_{t,\tau}$ can be either equal to $b > 0$ or to 0.[5] Corresponding to these two possibilities we will say that the payoff for a particular DM for a particular period is "biased" or "unbiased" respectively.

Differences in bias between the two players are the sources of preference misalignments between the two DMs. Hence, they are the key drivers of all equilibrium decisions in the model. We consider two canonical cases.

The first we call *Ideological Bias*. In this case we set $b_{1,1} = b_{1,2} = b > 0$, and $b_{2,1} = b_{2,2} = 0$. In other words, DM1 is biased in *both* periods, whereas DM2 is biased in *neither* of the two periods. DM1 systematically favors larger values of a_1 and a_2 relative to DM2.[6] Our terminology is motivated by the fact that the ideological bias model applies naturally to political environments, where an outgoing political party has systematically different policy preferences than its newly elected rival.

We term the second canonical case *Agency Bias*. In this case $b_{1,1} = b_{2,2} = b > 0$ and $b_{1,2} = b_{2,1} = 0$. In this case the preference misalignment is driven by the role that each DM has in each period: active policy maker or passive stakeholder. Each 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." 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 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

[5] Given the quadratic form of payoffs in (1), picking one of the two values as 0 is just a normalization of course.

[6] The actual *direction* of the preference misalignment — DM1's ideal point is larger than DM2's — is clearly inessential.

principal, or, in our model, DM2 in period 1 and DM1 in period 2).

The agency bias case captures the conflict within firms whereby a manager's preferences are misaligned with his or her successor or predecessor. The current manager, for instance, 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.

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

	$\tau = 1$	$\tau = 2$
DM1	b	b
DM2	0	0

Ideological Bias

	$\tau = 1$	$\tau = 2$
DM1	b	0
DM2	0	b

Agency Bias

(2)

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

Before proceeding further we illustrate with a simple application the case of Ideological Bias. Consider a government's problem of choosing its carbon emissions policy. The government faces turnover in its leadership. Specifically, two political parties take turns occupying the leadership position. In the first period, the "Carbon Party" occupies power, only to be replaced in the second period by the "Green Party." When in power, each party must choose a carbon emissions policy. The policy choice a_t may be interpreted as an attempt by the relevant political leader to hit a desired emissions target. The parameter θ then embodies the existing state of climate science. For any given θ, the Carbon Party that occupies power first prefers the emission level $\theta + b$ while then the Green Party prefers an emissions level that targets θ.[7] We return to political turnover in Section 4 below.

In the next 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.

[7]This carbon emissions example is one concrete instance of the "political turnover" model. Others include inflation targeting (θ is the inflation target) and anti-terrorism policy (θ is the threat level).

3.2. Information, Auditing, and Disclosure

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 $\mathbb{R}$. We will sometimes refer to DM1 as the "informed agent" and DM2 as the "uninformed agent."

In the carbon emissions policy example, both the state of climate science θ and the emissions level a_1 would be difficult to measure and, hence, are not automatically observed by anyone other than the party in power.

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.

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. We will refer to options to investigate or verify as a *policy audit* and a *state verification* respectively.

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) she can still choose not to audit, and therefore rely only on the message (we denote this choice by NA for "no audit"), (ii) she can undertake a costless audit to verify the policy choice a_1 of DM1 (we denote this choice by PA, for "policy audit"), or (iii) she can 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 ad PA are "perfect." State Verification fully reveals the state θ, while a Policy Audits fully reveals the first period policy choice a_1.

We let the choice across these three possibilities be denoted by d so that

$$d \in \{\text{NA}, \ \text{PA}, \ \text{SV}\} \tag{3}$$

Other things equal, DM2 would clearly (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 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.

In order to create reasonable trade off between SV and PA, we assume that Policy Audits are costless while State Verification carries a small cost.[8] The cost will be so small, in fact, as to be lexicographic; this of course implies that State Verification is chosen unless DM1's action a_t fully reveals the state θ. In this case a Policy Audit is (lexicographically) cheaper for DM2 and since it is equally informative it will be chosen over SV.

To summarize, temporally distinct decision makers (managers, political leaders, etc.) have conflicting preferences about the ideal policy each period. Preferences are influenced by an underlying state that is privately observed by the initial decision maker. The policy of DM1 is not (immediately) observed by DM2. DM2 can choose whether or not to forgo the Policy Audit and State Verification options and commit to observing only DM1's cheap talk message m. DM1 chooses his message and action after observing DM2's commitment choice. If DM2 does not commit, she can later choose between still only observing the message m (denoted NA), observing DM1's action (denoted PA) and verifying the actual value of θ (denoted SV).

The timing is depicted schematically below.

[8]For our purposes, it matters only that PA should be is less costly than SV. The main results focus on the surprising outcome that DM2 in equilibrium may commit to disclosure and foregoes PA and SV. So adding a cost to PA will only reinforce this result.

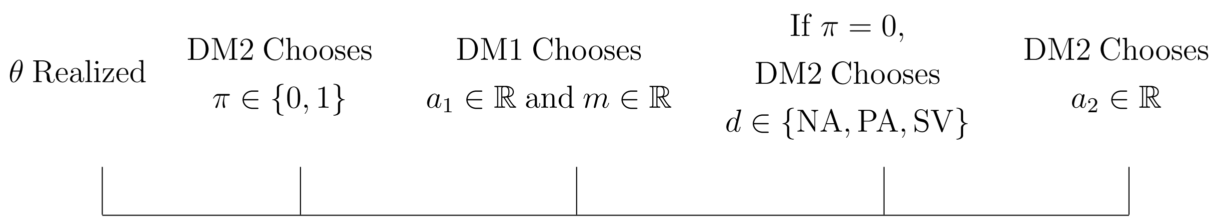

Figure 1

3.3. Equilibrium

We denote by z the actual observed outcome corresponding to the choice d at the discretionary stage. Specifically, z is the observed outcome when DM2 initially chooses *no commitment* ($\pi = 0$) and then chooses d. Note that corresponding to the three possible values of d as in (3) z will be equal to m, a_1 and θ respectively. If DM2 chooses $\pi = 1$ instead, of course she only observes m. Since all three possible observed variables take values in $\mathbb{R}$, a parsimonious way to describe what DM2 observes *overall*, taking both her choice of π and d into account is $\pi m + (1 - \pi)z$.

We let $a_1 = \alpha_1(\theta, \pi)$ denote the choice of policy by DM1 given his observation of θ and DM2's choice of π. Similarly, we let $m = \mu(\theta, \pi)$ denote DM1's choice of message m.

We also let $a_2 = \alpha_2(\pi\, m + (1 - \pi)z)$ denote DM2's choice of policy as a function of the actual value of the variable she observes given her choice of π and d.

Throughout the rest of the paper, by equilibrium we mean a Perfect Bayesian Equilibrium (PBE) of the game at hand.

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 the actual value of $\pi m + (1 - \pi)z$.*

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 the equilibrium has the feature that DM2 chooses $\pi = 0$, so that she does not commit to observing m only and has the option to conduct a policy audit or verify the state at a later stage. Then*

i) *It must be the case that $d = \text{PA}$ so that DM2 in fact chooses a policy audit when given the option of doing so.*

ii) *DM1's choice of $a_1 = \alpha_1(\theta, 0)$ of DM1 is fully revealing of the state θ.*

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.

Assume that in some equilibrium $\pi = 0$. When DM2 chooses d it is not possible that she will set $d = \text{NA}$ 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 wold not be fully revealing of the state. Hence DM2 could deviate to setting $d = \text{SV}$, observe θ and hence attain a higher payoff.

Could it then be that DM2 sets $d = \text{SV}$ when given the chance? 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 SV carries a lexicographic cost,[9] if the choice of a_1 is fully revealing of θ, DM2 would then deviate to choosing PA. This is just as informative as SV, but carries no cost.

Hence, the 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 the cost of choosing SV is lexicographic[10] 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 $d = \text{PA}$ 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.

[9]Note that the argument here does not hinge on the cost actually being lexicographic, but only "sufficiently small."

[10]The same remark we made in footnote 9 applies here.

From Proposition 1, the only candidates for equilibrium, in either the case of agency or ideological bias, are either for DM2 to either rely on the "words" (the case of $\pi = 1$) or the "actions" (the case of $\pi = 0$) of DM1.

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 fits well the case of two political parties who hold power in sequence but have a pre-set difference of opinion regarding the desirable values of a given policy variable.

Using (1) and (2) the payoffs to DM1 and DM2 in the case of ideological bias are respectively given by

$$\begin{aligned} V_1 &= -(a_1 - \theta - b)^2 - (a_2 - \theta - b)^2 \\ V_2 &= -(a_1 - \theta)^2 - (a_2 - \theta)^2 \end{aligned} \tag{4}$$

From Proposition 1 we know that only two alternatives are viable in equilibrium. Either DM2 chooses to *commit* to observing only m and 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 ($d =$ PA) to observe a_1. 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.

Especially 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.

The main result for the case of Ideological Bias tells us that this intuition turns out to be false for "large enough" values of b. 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 $\alpha_1(\theta, 0) = \theta + 2b$.*

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, whenever it exists, we will refer to this 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 $\overline{b} > 0$ such that*

i) *Whenever $b > \overline{b}$ in any equilibrium DM2 selects $\pi = 1$. In other words 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) *Whenever $b \leq \overline{b}$ the model has 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 $a_1 = \alpha_1(\theta, 0) = \theta + 2b$.*

Obviously, the more surprising part of Proposition 3 is part i. In the case of ideological bias, when b is sufficiently large, words speak louder than actions. DM2 commits to observe only m in this case. The full proof of Proposition 3 is in the Appendix. Here we develop an intuition for the result in part i, and embed it in the political ideology interpretation that we favor. We then comment on part ii.

The parameter θ represents DM1's information about the real state of climate change. The larger the θ the more acceptable it is to have larger emissions of CO_2. The first decision maker DM1 is the Carbon Party that we alluded to above. For any given level of θ the Carbon Party's ideal level of emissions is $\theta + b$ both in period $\tau = 1$ and $\tau = 2$. The second decision maker DM2 is the Green Party, that we alluded to above. For any given level of θ they favor emissions that equal to θ in both periods.

If b is large enough, the Green Party, the "receiver" of information chooses to set $\pi = 1$ and so commits to observe m only. By doing so the Green Party looses 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 Green Party chooses $\pi = 1$ the Carbon Party knows it is choosing a_1 "in secret." Hence it will simply set $a_1 = \theta + b$. Hence, regardless of θ the first period payoff of the Green Party is $-b^2$ (see (4)). The Green Party chooses a_2 after observing m, and hence (using (4) again) its second period expected payoff contingent on m is $-Var(\theta|m)$.[11]

If the Green Party 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 the Green Party sets $\pi = 0$, the Carbon Party knows that its policy a_1 will be audited by the Green Party in the second period. Observing a_1, by Proposition 1 will give the Green Party *full information* about θ, who will therefore set $a_2 = \theta$, and hence give the Green Party a second period payoff of 0.

Knowing that the Green Party will use the information yielded by the Policy Audit to set $a_2 = \theta$ will induce the Carbon Party to "compensate" due to the marginal effect of its first period action on the choice of the Green Party, and set a value of a_1 that is above $\theta + b$. In fact we know from Proposition 2 that in the best case for the Green Party (and, for that matter, for the Carbon Party as well) a_1 will be set to equal $\theta + 2b$. Hence the first-period payoff to the Green Party 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 the Green Party is zero, this means that its overall expected payoff after setting $\pi = 0$ is at best $-4b^2$.

To sum up, the Green Party's expected payoff after setting $\pi = 1$ is bounded below by $-b^2 - Var(\theta)$, 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 the Green Party must prefer to set $\pi = 1$ and thus forgo any Policy Audit or State Verification and rely only on the cheap talk message of the Carbon Party m. 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 characterization of the equilibrium set in part ii of Proposition 3 is not quite as sharp as one would ideally like because of the underlying multiplicity of continuation equilibria both

[11] Clearly the actual value of $Var(\theta|m)$ depends on the entire message function $\mu(\cdot|0)$ chosen by the Carbon Party 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 the Carbon Party as balancing the marginal "tangible" payoff from a_1 (as it enters directly the Carbon Party's first period payoff) and its marginal effect (via full revelation) on the beliefs of the Green Party about θ and hence on a_2. This gives rise to a simple differential equation (see (A.1) in the Appendix) which underpins the linear equilibrium of Proposition 2.

following DM2 setting $\pi = 0$ and setting $\pi = 1$.

The statement of Proposition 3 is supported (see the Appendix for the details of the argument) by choosing a babbling equilibrium 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 clearly not a babbling equilibrium).

The answer is given formally in Proposition A.1 in the Appendix. In short, under a weak technical assumption on the distribution of θ (see Assumption A.1), if we restrict attention to equilibria that satisfy the forward induction requirement we have outlined, we find that there exists some $\underline{b} < \overline{b}$, such that whenever $b \leq \underline{b}$ then the *unique* surviving equilibrium is one in which DM2 chooses to set $\pi = 0$, hence retaining the option to Audit or to Verify, and then proceeds to Audit DM1's policy.

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 (2)) 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.

Using (1) and (2) the payoffs to DM1 and DM2 in the case of agency bias are respectively given by

$$\begin{aligned} V_1 &= -(a_1 - \theta - b)^2 \, - \, (a_2 - \theta)^2 \\ V_2 &= -(a_1 - \theta)^2 \, - \, (a_2 - \theta - b)^2 \end{aligned} \tag{5}$$

[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). See Definition A.2 for a formal statement of how we proceed here.

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 4. *Under Agency Bias Actions Speak Louder Than Words: Consider the case of Agency Bias. Then, for any level of b, there is an equilibrium 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 enjoys her globally optimal payoff of* 0.

iii) *The payoff to DM1 is no smaller than the one she gets in any other equilibrium in which DM2 sets $\pi = 0$.*

Note that combining Proposition 4 with Proposition 1 we know that in the equilibrium singled out by Proposition 4, DM2 proceeds to set $d =$ PA and so to audit DM1's policy and observe a_1 directly, thus gaining access to full knowledge of the actual value of θ.

Just as in the case of Proposition 3 the actual statement of the result is not quite as strong as one would like because of the underlying multiplicity of equilibria with $\pi = 0$. However, in this case too, a basic forward induction argument will rule out any equilibrium in which it is not the case that $\pi = 0$. In fact, even more sharply, a deviation to setting $\pi = 0$ from a putative equilibrium in which $\pi = 1$ will, by forward induction, be interpreted as a signal that he intends to play an equilibrium which he prefers to the putative one with $\pi = 1$ — namely the one that gives her the *globally* optimal payoff of 0.

Intuitively, the reason for the sharp difference in behavior when b is large between the Ideological Bias and the Agency Bias is not hard to explain.

In both cases, the action of 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 $\pi = 0$.

Just as in the case of Ideological Bias, 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. Under Agency Bias DM1 would like, if possible, to distort DM2's action downward via signaling, while the bias term pushes up DM1's action via its tangible value. With the particular payoff functions we are working with,[14] the two effects cancel out making DM1's choice equal to the one that is the overall favorite of DM2, namely $a_1 = \theta$.[15] 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

Three "robustness checks" are worth discussing. This first concerns discounting. The second concerns the role of quadratic payoffs. The third concerns DM2 having also some private information.

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 $\alpha_1(\theta, 0) = \theta + (1 + \delta)\ b$. The thresholds $\overline{b}$ and $\underline{b} < \overline{b}$ need to be recalculated for the discounting case, but their existence remains exactly as stated in Proposition 3 and Proposition A.1. Proposition 4 part i also continues to hold.

In short the overall picture framed by Propositions 3, A.1 and 4 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.

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.

[14] See Section 5 concerning the robustness of our results to the case of non-quadratic payoffs and to the case in which second-period payoffs are discounted.

[15] The differential equation embodying DM1's Incentive Compatibility constraint (see (A.13) in the Appendix) now yields $a_1 = \theta$ as one of the solutions. See footnote 12 above.

Of course, it is a legitimate question to ask how crucial their role is in the qualitative flavor of the results we obtain here. While we have no general results in this regard, we offer two remarks here that we believe clarify the issue, at least to some extent.

The first is that the qualitative overall picture of our result does depend in an obvious way on some characteristics of the preferences we have postulated. In the same way that the results change qualitatively between the case of Ideological Bias and Agency bias, one cannot possibly expect the overall picture to be unchanged across too wide a variety of preferences.

The second is that there is nothing in our results that is "knife-edge" in terms of the preferences we ascribe to the players. A simple scrutiny of our arguments reveals that all of them go through unchanged for preferences that are sufficiently close to the one we use. A formal statement would be cumbersome, and even more so a formal proof.[16] However, there is a sense in which our results continue to hold in an open neighborhood of the time-separable quadratic preferences we have used throughout. The time-separable quadratic-loss preferences we use here can be "perturbed" without invalidating our results. The details are clearly out of place in the present paper.

Lastly, it is natural to ask what happens if DM2 has some private information as well. Say that the sate evolves between period 1 and 2, and that DM2 has some private information about the change. In this case, it is straightforward to show that all the results we have presented above hold unchanged. The model is fully robust to at least some cases of bilateral asymmetric information.

6. Governments, Firms, Audits and Partisan Investigations

We have pointed out above that in broad outline the Ideological Bias case fits better the political economy scenario in which two parties successively hold power, while the agency case is more suited to representing a firm in which two CEOs head the operations in succession. Although this view is an obvious oversimplification,[17] our remarks here are predicated on the assumption that these two broad classifications capture some important part of the real distinctions between firms and polities.

Our results indicate that the ability to commit to a system of disclosure, forgoing the option to Audit or Verify, has little or no value in the case of a firm. This, in our view, tallies

[16] For a start we would need to pick a topology, or a suitably wide parameterized class of preferences, to make precise the notion of open neighborhood in the space of possible preferences.

[17] For instance, as we remarked above, office-motivated politicians — as opposed to the ideology-driven ones we have in mind here — may fit better the case of Agency Bias.

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.[18]

The picture changes considerably in the political sphere. In this case, our results indicate that the ability to commit to a disclosure regime — successfully forgoing the option to Audit or Verify — has value to the players.

Alas in many cases the ability to *commit* to a disclosure regime is virtually non-existent in the political sphere. Almost by definition, when a new party occupies power it has wide powers to renege on previous promises, if necessary by changing the law. Our 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.

In our view this is at least part of the reason why "audits" in the political sphere are often viewed as "partisan witch hunts"[19] where "all actors involved in accountability processes use a variety of strategies to argue their case and apportion blame" (Boin, McConnell, and Hart (2008)).

Indeed, the frequency of partisan investigations is an indication that commitments to avoid auditing/investigating are 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 results illustrate the destructive nature of such investigations in the case of ideologically-driven politicians. The question arises as to whether commitment could effectively be achieved in any political system, and if so, what form it would take. Clearly, a question worthy of future research.

[18] See Arter (2000).

[19] One instance of this is the recent climate science investigation in the U.S. The civil fraud investigation, initiated by the State of Virginia's Attorney General, a well known climate science skeptic, was recently thrown out in court and drew criticism for having been initiated for ideological reasons. See Nature (2010).

References

Anderlini, L., D. Gerardi, and R. Lagunoff (2012): "Communication and Learning," *Review of Economic Studies*, forthcoming.

Arter, D. (2000): "Management Auditing," *Quality Digest*, 73.

Austen-Smith, D., and J. Banks (2000): "Cheap Talk and Burned Money," *Journal of Economic Theory*, 91, 1–16.

Boin, A., A. McConnell, and P. Hart (2008): *Governing after Crisis: The Politics of Investigation, Accountability and Learning*. Cambridge: Cambridge University Press.

Carrillo, J., and T. Mariotti (2000): "Strategic Ignorance as a Self-Disciplining Device," *Review of Economic Studies*, 67, 529–44.

Crawford, V. P., and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431–1351.

Daughety, A., and J. Reinganum (2008): "Communicating Quality: A Unified Model of Disclosure and Signaling," *RAND Journal of Economics*, 73, 973–989.

——— (2010): "Public Goods, Social Pressure, and the Choice Between Privacy and Publicity," *American Economics Journal: Microeconomics*, 2, 191–221.

Farrell, J., and M. Rabin (1996): "Cheap Talk," *Journal of Economic Perspectives*, 10, 103–118.

Goldberg, A. (2010): "Been there, done that: Partisan investigations solve little," *Politico*.

Govindan, S., and R. Wilson (2009): "On Forward Induction," *Econometrica*, 77, 1–78.

Heim, R. (2011): "How to Win an Insider Trading Case," Meyers & Heim LLP. Unpublished, available at http://www.meyersandheim.com/insider_trading.html.

Karamychev, A., and D. Visser (2011): "An Optimal Signaling Equilibrium," Discussion paper, Erasmus University and the Tinbergen Institute.

Kartik, N. (2007): "A Note on Cheap Talk and Burned Money," *Journal of Economic Theory*, 136, 749–758.

——— (2009): "Strategic Communication with Lying Costs," *Review of Economics Studies*, 76, 1359–1395.

KRISHNA, V., AND J. MORGAN (2008): "Cheap Talk," in *The New Palgrave Dictionary of Economics*, ed. by L. Blume, and S. Durlauf, 2^{nd} edn. New York, NY: McMillan.

MAN, P. Y. (2012): "Forward Induction Equilibrium," *Games and Economic Behavior*, 75, 265–76.

NATURE, Editorial. (2010): "Science Subpoenaed," *Nature*, 465, 135–136, avaliable at http://www.nature.com/nature/journal/v465/n7295/full/465135b.html.

SOBEL, J. (2007): "Signaling Games," UCSD, mimeo.

SPECTOR, D. (2000): "Rational Debate and One-Dimensional Conflict," *Quarterly Journal of Economics*, 115, 181–200.

SPENCE, A. M. (1973): "Job Market Signaling," *Quarterly Journal of Economics*, 90, 225–243.

VAN DAMME, E. (1989): "Stable Equilibria and Forward Induction," *Journal of Economic Theory*, 48, 476–496.

WATERGATE HEARINGS (1973): *Hearings Before The Select Committee on Presidential Campaign Activities Of The Unites States Senate. Ninety-Third Congress, First Session. Watergate and Related Activities. Book 4.* Washington, DC: United States Government Printing Office.

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 $\alpha_1(\theta, 0)$ 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 $a_1 = \alpha_1^*(\theta, 0)$ must solve

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

Assuming α_2^* is differentiable,[20] we can write the first order condition as

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

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

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

$$[\alpha_1^*(\theta, 0) - \theta - b] \; \alpha_1^{*\prime}(\theta, 0) \; - \; b \; = \; 0 \tag{A.1}$$

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

This clearly suffices to prove part ii of Proposition 2. For future reference, at this point we notice that $\alpha_1^*(\theta, 0) = \theta + 2b$ 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 ii of Proposition 1, it is sufficient to show that the payoffs associated with the linear equilibrium in which $\pi = 0$ and $\alpha_1^*(\theta, 0) = \theta + 2b$ Pareto-dominate those of any other fully revealing equilibrium in which DM2 sets $\pi = 0$.

Consider again the differential equation in (A.1), which we know any fully revealing equilibrium to be considered here must satisfy (almost everywhere).

Define

$$g(\theta) \; = \; \alpha_1(\theta, 0) - \theta - b \tag{A.2}$$

and notice

$$g'(\theta) \; = \; \alpha_1^{*\prime}(\theta, 0) - 1 \tag{A.3}$$

Hence, we can rewrite (A.1) as

$$g(\theta)\,[1 + g'(\theta)] = b \tag{A.4}$$

[20] 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 clearly not a viable option. However, it suffices to notice that 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.

Clearly, $g^*(\theta) = b$ is one solution to (A.4), giving rise to the linear equilibrium strategy $\alpha_1^*(\theta, 0) = \theta + 2b$ that we identified above.

In general, however, there are infinitely many solutions to (A.4), corresponding to infinitely many fully revealing equilibria.[21]

Once we have a g that solves (A.4), for any given θ the payoffs to the decision makers in the corresponding equilibrium, can be respectively written as

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

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 (A.4) and $g \neq g^*$, then $g(\theta) \geq b$ for every θ. This will clearly 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 (A.4). 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 (A.4) 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 (A.4) 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 $\lim_{\theta \to -\infty} \tilde{g}'(\theta) = 0$. Since $\tilde{g}(\theta) < b$ for all $\theta < \hat{\theta}$, this implies that for some sufficiently small $\overline{\theta}$ it must be that

$$\tilde{g}(\overline{\theta})[1 + \tilde{g}'(\overline{\theta})] \; < \; b \tag{A.5}$$

which directly implies that $\tilde{g}$ cannot be a solution to (A.4).

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 $\lim_{\theta \to -\infty} \tilde{g}'(\theta) = 0$. Since $\tilde{g}(\theta) < 0$ for all $\theta < \hat{\theta}$, this implies that for some sufficiently small $\overline{\theta}$ it must be that

$$\tilde{g}(\overline{\theta})[1 + \tilde{g}'(\overline{\theta})] \; \leq \; 0 \tag{A.6}$$

which directly implies that $\tilde{g}$ cannot be a solution to (A.4). ■

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 $-b^2$ $-Var(\theta)$.*[22]

[21] 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 (A.4). Suppose $\overline{g}$ is a solution to (A.4). Fix Δ and define $\tilde{g}$ by letting $\tilde{g}(\theta) = \overline{g}(\theta + \Delta)$ for every θ. Then $\tilde{g}$ is also a solution.

[22] Here $E(\theta)$ denotes the expected value of θ.

Proof: Given that DM2 chooses $\pi = 1$, DM1's choice of a_1 will not be observed by DM2. Hence, from (4) 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 (4), ignoring m and setting $a_2 = E(\theta)$, it is clear that DM2 can achieve a second-period expected payoff of $Var(\theta)$. ■

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 (4), 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 $-4b^2$. ■

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

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 $-b^2 - Var(\theta)$. 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 $\overline{m}$ 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 $-b^2 - Var(\theta)$. Since $b \leq \bar{b}$ this is not a profitable deviation and hence the proof of part ii is now complete. ■

Definition A.1. *Auxiliary Games: Consider the general model defined in Section 3. Now consider two auxiliary games. One obtained form the original game, but in which DM2 has no choice of π and this is simply exogenously set to be equal to 1. We call this the* exogenous disclosure *game. Further, consider another game obtained from the original one in which DM2 chooses neither π nor d, and π is exogenously set to be equal to 0 and d is exogenously set to be equal to PA. We call this second one the* exogenous audit *game.*

Remark A.1: *Consider an equilibrium of the general Model defined in Section 3. Note that since Definition 1 restricts attention to Perfect Bayesian Equilibria, we know that the players' equilibrium strategies conditional on $\pi = 1$ constitute a Perfect Bayesian Equilibrium of the exogenous disclosure game of Definition A.1. Note also that since Definition 1 restricts attention to Perfect Bayesian Equilibria, using also Proposition 1, we know that the players's strategies conditional on $\pi = 1$ constitute a Perfect Bayesian Equilibrium of the exogenous audit game of Definition A.1.*

Definition A.2. *Forward Induction Proof: Consider an equilibrium of the general Model defined in Section 3. We say that this equilibrium is* forward induction proof *(henceforth a FIPE) if and only if it satisfies the following two requirements.*

The players' strategies conditional on $\pi = 1$ yield a payoff to DM2 that is no smaller than the payoff that DM2 obtains in any equilibrium of the exogenous disclosure game.

The players' strategies conditional on $\pi = 0$ yield a payoff to DM2 that is no smaller than the payoff that DM2 obtains in any equilibrium of the exogenous audit game.

Before moving on to the statement and proof of Proposition A.1 we state an assumption (needed for the proposition), a definition and two Lemmas that will be used in the proof.

Assumption A.1. *Constant Density Over $[\underline{\theta}, \overline{\theta}]$: The strictly positive continuous density f from which θ is drawn is constant on some non-degenerate interval $[\underline{\theta}, \overline{\theta}]$. Henceforth we let φ denote the the probability that $\theta \in [\underline{\theta}, \overline{\theta}]$.*

Definition A.3. *Bounded Partitional: An equilibrium of the exogenous disclosure game 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 equilibria of the exogenous audit game 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. ■

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

Consider any sequence of associated equilibria of the exogenous disclosure game, and let V_{2k}^* be the sequence of associated payoffs for DM2.

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

Then, there exists a $\overline{k}$ such that for every $k \geq \overline{k}$, 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 $[\underline{\theta}, \overline{\theta}]$. Suppose only one message is sent, say $\hat{m}$. Conditional on $\theta \in [\underline{\theta}, \overline{\theta}]$, DM2's payoff is then bounded above by $-Var(\theta | \theta \in [\underline{\theta}, \overline{\theta}])$. Suppose next that only two message are sent, say $\hat{m}$ and $\hat{\hat{m}}$. Then the payoff to DM2 is bounded above by $-Var(\theta | \theta \in [\underline{\theta}, \overline{\theta}] \,\text{and}\, m = \hat{m}) - Var(\theta | \theta \in [\underline{\theta}, \overline{\theta}] \,\text{and}\, m = \hat{\hat{m}})$. 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 $[\underline{\theta}, \overline{\theta}]$ the second period payoff to DM2 is bounded above by $-\ell\ \varphi$. This is because DM2's payoff conditional on $\theta \notin [\underline{\theta}, \overline{\theta}]$ is at most 0. Since in any equilibrium of the exogenous disclosure game 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 $[\underline{\theta}, \overline{\theta}]$ in any equilibrium of the exogenous disclosure game 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. ■

Proposition A.1*: Consider the Ideological Bias case and let Assumption A.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 $d = PA$.*

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

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 $[\underline{\theta}, \overline{\theta}]$.

Our next endeavor is to compute an upper bound to DM2's payoff contingent on $\pi = 0$ — or equivalently in the exogenous disclosure game.

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 $\overline{\theta}$. Let such bound be denoted by $\overline{d}_k$. Without loss of generality (by taking subsequences if necessary) we assume that both sequences $\{\underline{d}_k\}_{k=0}^{\infty}$ and $\{\overline{d}_k\}_{k=0}^{\infty}$ converge as $k \to \infty$.

Notice next that it must be the case that

$$\lim_{k\to\infty} (\overline{d}_k - \underline{d}_k) \;=\; \eta \;>\; 0 \tag{A.7}$$

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 $\theta \in [\underline{d}_k, \overline{d}^k]$ and $\lim_{k\to\infty} \phi_k = \phi$. Note that by Assumption A.1 it must be that $\phi = \eta\, \varphi/(\overline{\theta} - \underline{\theta})$ and hence $\phi > 0$.

Notice that at this point we know that contingent on $\theta \in [\underline{d}_k, \overline{d}^k]$, the exogenous disclosure game is exactly the model in CS in the special case of a uniform density.[23]

Let $U(\underline{d}_k, \overline{d}^k, b_k)$ be DM2's second period payoff under the most informative (the one that gives DM2 the highest payoff), contingent on $\theta \in [\underline{d}_k, \overline{d}^k]$. Overall, DM2's equilibrium payoff in the exogenous disclosure game is then bounded above by

$$-b_k^2 + \phi_k U(\underline{d}_k, \overline{d}_k; b_k) \tag{A.8}$$

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

Our contradiction hypothesis and (A.8) imply that there exists $n \in \mathbb{N}$ such that for every $k \geq n$,

$$-b_k^2 + \frac{\phi}{2} U(\underline{d}_k, \overline{d}_k, b_k) \geq -4b_k^2 \tag{A.9}$$

which implies

$$U(\underline{d}_k, \overline{d}_k; b_k) \geq -\frac{6b_k^2}{\phi} \tag{A.10}$$

Directly from CS in the constant density case[24] we get that

$$U(\underline{d}_k, \overline{d}_k, b_k) \;=\; -A_k b_k^2 - B_k \tag{A.11}$$

---

[23]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.

[24]See footnote 23 above.

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

$$U(\underline{d}_k, \overline{d}_k, b_k) \ = \ -A_k b_k^2 - B_k \ < \ -\frac{6b_k^2}{\phi} \tag{A.12}$$

which contradicts inequality (A.10) and hence concludes the proof ■

We continue with a preliminary result that will be used to prove Proposition 4.

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 $\alpha_1(\theta) = \theta$.*

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.*

Proof: [25] Consider 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 $\alpha_1(\theta, 0)$ 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 $a_1 = \alpha_1^*(\theta, 0)$ must solve

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

Assuming α_2^* is differentiable,[26] we can write the first order condition as

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

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

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

$$[\alpha_1^*(\theta, 0) - \theta - b] \, \alpha_1^{*\prime}(\theta, 0) \ + \ b \ = \ 0 \tag{A.13}$$

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

This clearly 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 $\alpha_1^*(\theta, 0) = \theta$ Pareto-dominate those of any other fully revealing equilibrium in which DM2 sets $\pi = 0$.

[25] 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.

[26] See footnote 20 above.

Consider again the differential equation in (A.13), which we know any fully revealing equilibrium to be considered here must satisfy (almost everywhere), and define

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

and notice

$$g'(\theta) \;=\; \alpha_1^{*\prime}(\theta, 0) - 1 \tag{A.14}$$

Hence, we can rewrite (A.13) as

$$g(\theta)\,[1 + g'(\theta)] \;=\; -b \tag{A.15}$$

Clearly, $g^*(\theta) = -b$ is one solution to (A.15), giving rise to the linear equilibrium strategy $\alpha_1^*(\theta, 0) = \theta$ that we identified above.

In general, however, there are infinitely many solutions to (A.15), corresponding to infinitely many fully revealing equilibria.[27]

Once we have a g that solves (A.15), for any given θ the payoffs to the decision makers in the corresponding equilibrium, can be respectively written as

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

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 (A.15) and $g \neq g^*$, then $g(\theta) \leq -b$ for every θ. This will clearly 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 (A.15). 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 (A.15) 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 (A.15) 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 $\lim_{\theta\to\infty} \tilde{g}'(\theta) = 0$. Since $\tilde{g}(\theta) > -b$ for all $\theta > \hat{\theta}$, this implies that for some sufficiently large $\overline{\theta}$ it must be that

$$\tilde{g}(\overline{\theta})[1 + \tilde{g}'(\overline{\theta})] \;>\; -b \tag{A.16}$$

which directly implies that $\tilde{g}$ cannot be a solution to (A.15).

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 $\lim_{\theta\to\infty} \tilde{g}'(\theta) = 0$. Since $\tilde{g}(\theta) > 0$ for all $\theta > \hat{\theta}$, this implies that for some sufficiently large $\overline{\theta}$

[27] See footnote 21 above.

it must be that

$$\tilde{g}(\overline{\theta})[1+\tilde{g}'(\overline{\theta})] \;\geq\; 0 \tag{A.17}$$

which directly implies that $\tilde{g}$ cannot be a solution to (A.15). ■

Proof of Proposition 4: We simply show that DM2 setting $\pi = 0$ followed by the linear equilibrium identified in Lemma A.5 constitutes an overall equilibrium of the model.

The claim is in fact rather obvious. Since in this equilibrium $\alpha_1^*(\theta, 0) = \theta$, it is clear that DM2 achieves her globally optimal payoff of 0. Hence it cannot be the case that DM2 strictly prefers to deviate from choosing $\pi = 0$. ■