American Economic Review 100 (June 2010): 1136–1162
http://www.aeaweb.org/articles.php?doi=10.1257/aer.100.3.1136

Peddling Influence through Intermediaries

By Wei Li*

A sender may communicate with a decision maker through intermediaries. In this model, an objective sender and intermediary pass on information truthfully, while biased ones favor a particular agenda but also have reputational concerns. I show that the biased sender and the biased intermediary's reporting truthfulness are strategic complements. The biased sender is less likely to use an intermediary than an objective sender if his reputational concerns are low, but more likely to do so if his reputational concerns are moderate. Moreover, the biased sender may be more likely to use an intermediary perceived to be more biased. (*JEL* D82, D83)

A government intent on pushing a particular agenda or selling a policy may convey the relevant information to the public directly. However, doing so may be risky, especially if the agenda is unsupported by later evidence or the policy turns out wrong. The government may also convey its information to the media, both traditional and online, under condition of anonymity ("background briefing" only).[1] The media then chooses what to tell the public. Such practices are common; for instance, prewar intelligence on Iraq was intentionally leaked to news media. The 2007 trial and conviction of I. Lewis Libby Jr. unfolded in a manner that suggests senior administrators had disclosed classified information to reporters for political purposes.[2] What are the pros and cons of influencing public opinion through intermediaries for the government?

This paper develops a model of communication through strategic intermediaries. A partially informed sender—the government—sends a message to an intermediary, a media outlet, which then sends a message to the uninformed decision maker, the public. The public takes an action based on what it hears but eventually observes the true state. The government and the intermediary can each be objective or biased: an objective agent is assumed to pass on information truthfully, but a biased one wants to push a particular agenda *and* to appear objective. A biased government must balance two opposing considerations. On one hand, the intermediary reduces the government's reputation cost of releasing inaccurate information, because the public may think that it had been misled by the messenger if the later observed true state contradicts the information from the intermediary. This *blame sharing effect* makes it more attractive for the government to use the intermediary. On the other hand, the intermediary also reduces the

* Department of Economics, University of California at Riverside, and University of British Columbia, Buchanan Tower 1027, Vancouver, BC V6T1Z1, Canada (e-mail: weiliubc@interchange.ubc.ca). I thank Glenn Ellison, Botond Kőzegi, Edward E. Schlee, and especially Li Hao for insightful suggestions and detailed comments. I have also benefited from the thoughts of Mathias Dewatripont, Tracy Lewis, R. Preston McAfee, Marco Ottaviani, Bryan A. Richards, Jean Tirole, three anonymous referees, and the seminar participants at MIT, Kellogg School of Management, UCR, ASU, USC, UCSB, Academia Sinica, Taiwan, Johns Hopkins University, HKUST, Hong Kong University, UBC, the 9th Southwest Economic Theory conference, the 7th Canadian Economic Theory Conference, the 2007 Summer Meeting of the North American Econometric Society, and the Duke/Northwestern/Texas IO conference.

[1] Anonymity is widely granted in the news media. For instance, in the first week of April 2005, 47 percent of all A-section articles published in the *New York Times* used anonymous sources, 46 percent of which were identified as "officials" or "aides" only. See "Briefers and Leakers and the Newspapers Who Enable Them," Daniel Okrent, *New York Times*, May 8, 2005.

[2] "Libby Guilty of Lying in C.I.A. Leak Case," A. Neil Lewis, *New York Times*, March 7, 2007.

effectiveness of any message the government uses to push its agenda because, having no information of its own, the intermediary introduces possible distortions without adding to the message's accuracy. This *credibility reducing effect* makes it less attractive for the government to use the intermediary.

The net effect of using an intermediary on the government may seem ambiguous. The government, however, is shown to report *less* truthfully using an intermediary who may have the same bias. This is a direct consequence of the first insight emerging from this model: because the blame sharing effect outweighs the credibility reducing effect, the government and the intermediary's truth telling probabilities are strategic complements. If one exogenously gains stronger incentives to report more truthfully, perhaps due to a higher reputation cost of lying, so does the other due to the strategic complementarity. Since for the government, communicating directly is equivalent to using a truthful intermediary, using a possibly biased intermediary who may distort his message enables the government to report less truthfully as well.[3] The blame sharing effect outweighs the credibility reducing effect because of a crucial difference in the information available to the decision maker when she takes her action based on the media's report and when she evaluates the government and the media *after* observing the true state. Consider the critical event when the intermediary's message is associated with bias, and that message contradicts the true state. When the decision maker takes an action, she allows for the possibility that the intermediary's message results from objective agents who pass on information truthfully. When the decision maker evaluates the agents' objectivity, however, she has observed the true state. Because the intermediary's message is wrong, she assigns a higher probability to the event that the intermediary has distorted the information from the government. Hence the intermediary shares the government's blame more than it reduces the credibility of the government's message.

The public, then, should evaluate anything learned from an intermediary cautiously: the intermediary may not only introduce distortions of his own, but also worsen the government's incentives to report truthfully. Because the truth telling incentives of the biased government and the biased intermediary move in the same direction, the public can reduce the information loss even if the policy measures at its disposal cannot directly reach everyone. Policies such as stricter enforcement of disclosure laws or higher standards for granting anonymity increase the reputation cost of the intermediary to lie, and in turn make the government more truthful.

If the government with an agenda may choose to communicate directly or to use an intermediary before receiving his private information, the channel chosen becomes a signal of how objective he is, and the public interprets the message it receives accordingly. The government's channel choice hinges on how important, given its characteristics, the blame sharing effect from using the intermediary is relative to the credibility reducing effect. The second insight from this model is that a biased sender is more likely to choose direct communication than an objective one if his reputational concerns are sufficiently low. In this case, even though the biased sender pays a higher reputation cost both because direct communication becomes a negative signal of his objectivity, and because he has no intermediary to share the blame with, the reputation cost is strictly smaller than the gain in message credibility from direct communication. In contrast, he is more likely to use an intermediary if he has moderate reputational concerns. In this case, the blame sharing effect becomes more important for the biased sender than the loss in message credibility, despite the fact that indirect communication becomes a negative signal of his objectivity. Further, when both the biased sender and the biased intermediary have low reputational concerns, the probability that a biased sender chooses direct communication may increase in the

[3] Throughout this paper, the sender and the intermediary are male and the decision maker is female.

perceived objectivity of the intermediary. The reason is that an intermediary of sterling objectivity shares so little blame that it is not worth the loss in credibility for the biased sender to use him.

In the main model, the sender and the intermediary, if biased, want to push the same agenda. The analysis can be easily adapted to alternative types of intermediaries. In particular, this paper shows that a biased sender faces the same tradeoff between credibility reducing effect and blame sharing effect if the intermediary has an opposite bias; or if the intermediary simply adds noise to the communication process. Moreover, similar to the main model, the net effect depends on whether, given the possible distortions of the intermediary, the decision maker is more or less likely to think, after observing the true state, that the intermediary has distorted the sender's message. If the intermediary may have an opposite bias, for instance, the credibility reducing effect dominates and the biased sender and the biased intermediary's truth telling incentives are strategic substitutes. If the intermediary is unbiased but may send the opposite of the sender's message with a small probability (noise), either effect may dominate because noise always makes the sender's message less credible, but its effect on the sender's reputation cost varies. For instance if the biased sender is sufficiently concerned about his reputation, then he reports less truthfully with a small amount of noise than without. Because the biased sender reports very truthfully without noise, the decision maker mainly attributes a wrong message to noise.

In terms of the setup of the model without intermediaries, this paper is related to Joel Sobel (1985) and Roland Bénabou and Guy Laroque (1992). Sobel (1985) considers a model in which the objective type reports honestly, but the biased type needs to appear credible in order to manipulate the decision maker through possibly distorted messages. Stephen Morris (2001) endogenizes the role of the objective type so that an objective sender also faces reputational concerns. He shows that there may exist a "politically correct" equilibrium in which the message associated with the bias is avoided by an objective sender sufficiently concerned about his reputation. The objective agents in the current paper are not strategic, but their truthful reporting strategies within a communication channel can be supported in equilibrium of a model akin to Sobel (1985), while their channel choice behavior can be endogenized in a model akin to Morris (2001).

More recently, Ming Li and Tymofiy Mylovanov (2008) consider a model in which a strategic expert may follow the recommendation from a third party in exchange for a fee, or incur a cost to gather his own information. If captured by the third party, the expert is an intermediary as in the current paper, and thus faces a similar tradeoff between agenda pushing and reputational concerns; however, the sender (the third party) in their model is not strategic. Because in this paper both the sender and the intermediary may be strategic, it is possible to study their strategic interactions as well as the sender's choice of communication channels.

Following the seminal work of Vincent P. Crawford and Joel Sobel (1982), many have studied the incentives of a biased sender who aims to influence a receiver by manipulating the information he sends (David Austen-Smith 1990; Mathias Dewatripont and Jean Tirole 1999; Judith Chevalier and Glenn Ellison 1999; Morris 2001; Andrea Prat 2005; Marco Ottaviani and Peter Sorensen 2006). In these models, the informed sender always communicates directly with the decision maker. This paper instead gives conditions under which the sender may prefer indirect to direct communication—a step toward explaining the widespread use of communicating through intermediaries. Several recent papers also study the role of nonstrategic intermediaries by extending the Crawford and Sobel (1982) framework to more general communication protocols (Vijay Krishna and John Morgan 2004; Andreas Blume, Oliver Board, and Kohei Kawamura 2007; Maria Goltsman, Johannes Horner, Gregory Pavlov, and Francesco Squintani 2009). In particular, Blume, Board, and Kawamura (2007) show that adding noise to the communication process may enable more information to be transmitted than is possible in Crawford and Sobel (1982), partly because the noise dampens the receiver's response to any message from the sender's point of view, and thus reduces the sender's incentive to distort his signal. In the current paper, noise reduces the sender's message credibility in a similar fashion, but noise may also reduce the sender's reputation cost to such an extent that the sender lies more, not less, often.

Section I sets up the indirect communication model. Section II characterizes the equilibrium of the indirect communication game after demonstrating the strategic complementarity between the biased sender and the biased intermediary's truth telling incentives. Section III studies, ex ante, whether a biased sender chooses to use an intermediary or to communicate directly. Section IV extends the analysis to allow for other types of intermediaries. Section V discusses key assumptions and concludes. All the proofs are collected in the Appendix and the Web Appendix.

I. Indirect Communication: Model

There are two players in this game, agent A and B, and there is a decision maker C. The state of the world is binary: $\eta \in \{0, 1\}$. Each state occurs with equal probability. This game consists of three stages: information transmission, decision making, and evaluation. In the information transmission stage, agent A observes a private signal $s_A \in \{0, 1\}$, which is equal to the true state with probability $p_A > 0.5$; otherwise it is wrong. He sends a message $m_A \in \{0, 1\}$ to an intermediary, agent B, who has no information of his own. Agent B then sends a message $m_B \in \{0, 1\}$ to the decision maker. In the decision making stage, C chooses an action $a \in \mathbb{R}$ given message m_B. In the evaluation stage, C first observes the true state η and then forms posterior beliefs about the type of agent j, $j = A, B$, to be described next. In all three stages, agent B and decision maker C observe only the message sent directly to him (her).

The decision maker's payoff is represented by the quadratic loss function $-(a - \eta)^2$. Her optimal action is thus to choose a equal to the probability she attaches to $\eta = 1$. Agent j may be either objective (type o) or biased (type b). Each agent's type is independently drawn from $\{o, b\}$: $\Pr(j = o) = \theta_j$, $\Pr(j = b) = 1 - \theta_j$, with θ_j referred to as j's prior objectivity. An objective agent is assumed to report his information honestly. A biased agent has an agenda: he always wants action $a = 1$ taken, regardless of the true state; but he also wants to be perceived as objective. Let C's posterior belief of agent j's being objective be π_j, which is formed at the evaluation stage and referred to as j's posterior objectivity. Biased A, B's payoffs are respectively:

$$U_A = a + \alpha\pi_A \quad \text{and} \quad U_B = a + \beta\pi_B.$$

The first part of biased j's payoff function is C's action. The higher is C's action, the better off a biased agent is. The second part is a reduced form formulation representing a biased agent's reputational payoff, which is assumed to be linear in his respective posterior objectivity to reflect the idea that an agent perceived to be less objective will lose influence in the future.[4] Parameters $\alpha, \beta \in [0, \infty)$ are the weights biased A and B attach to their reputation. The biased agents may care about reputation for various reasons. For instance, biased A may be concerned about the chances of being reelected if he is perceived as biased. Biased B may care about the future subscriptions to its news services—his information has little value to the readers if he is perceived to have a strong bias. The weights α, β also have two alternative, and economically relevant interpretations. First, the lower is α and β, the more biased A and B care about pushing their agenda now than later. Second, α may vary in situations where A is among several agents who may send

[4] The linear functional form is used in many existing papers (David S. Scharfstein and Jeremy C. Stein 1990; Canice Prendergast and Lars Stole 1996; Ottaviani and Sorensen 2006). In general, however, the agents' reputational payoffs are determined by C's decision problem in the future and can be either linear or convex. For example, in a two stage game, Morris (2001) shows that a biased agent's reputational payoff is convex in his posterior objectivity if he cares about influencing the decision maker's future action; and Wei Li (2007) shows that an expert's reputational payoff is piecewise linear and convex in his perceived talent if his future wage is his expected value of information.

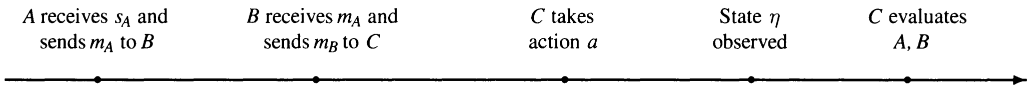

Figure 1. Timeline of the Indirect Communication Game

a message to B, in which case the decision maker assigns a smaller probability to any particular A being the source of B's information. Thus biased A's reputation is less affected by B's message, which amounts to his having lower reputational concerns.

In this game, a strategy of biased A consists of two probabilities of reporting truthfully, one for each signal $s_A = 0$ and $s_A = 1$. Analogously, a strategy of biased B is a probability of reporting truthfully for each message $m_A = 0$ and $m_A = 1$. This paper looks for perfect Bayesian equilibrium (PBE), in which given strategies of biased A and B, C's action a at the decision-making stage maximizes her (expected) payoff, and her posterior belief at the evaluation stage π_j satisfies Bayes's rule. The indirect communication game is summarized in Figure 1.

One key assumption of this model is that the objective type is nonstrategic and reports honestly. Honesty here is interpreted either as an institutional goal or a behavioral trait, similar to Sobel (1985) and Bénabou and Laroque (1992). Some media and nonprofit organizations may adhere to an ethical standard of informing the public only in an impartial way; people may simply prefer behaving honestly, as suggested by psychological experiments (John H. Evans, R. Lynn Hannan, Ranjani Krishnan, and Donald V. Moser 2001).[5] This assumption is further discussed in Section VA and relaxed in Web Appendix B.

Agent B is a pure intermediary and has no information of his own in this model. This reflects that A's information is of an exclusive nature. In the opening example, agent A has classified information unavailable to the media or the public. The main results of this paper hold qualitatively if B observes a sufficiently uninformative signal. The case where the intermediary is a well informed expert in the market for credence goods is discussed further in Section VB.

Before turning to the analysis, it is useful to keep in mind two possible applications of this model. In the first application, decision maker C represents voters who need to choose whether to vote for a candidate, agent A, in an upcoming election. Her optimal decision depends on the state of the world, which is whether A's opponent is involved in a scandal. Agent A may have evidence against the opponent and want to inform the voters ("objective type"); or he just wants to discredit the opponent ("biased type"). Biased A has reputational concerns: the voters may punish him in future elections if he is perceived as biased. Agent A may run a direct advertisement attacking the opponent. Under the Bipartisan Campaign Reform Act (BCRA) enacted in 2002, he must disclose his identity.[6] Alternatively, he may convey the information to intermediary B, a political action committee (PAC) or an activist group who may choose what to tell the voters. Such groups, for example the 527 organizations—tax exempt organizations that engage in political advocacy—are not subject to the same disclosure rules.

In the second application, decision maker C represents investors who need to make buy/sell decisions depending on the state of the world, which is whether a company is performing poorly. Agent A, a market insider, may have information about the company's poor performance

[5] For instance, BBC's editorial guideline states that "We will be objective and evenhanded in our approach to a subject. We will provide professional judgments where appropriate, but we will never promote a particular view on controversial matters of public policy or political or industrial controversy."

[6] Political candidates for federal offices need to comply with the "stand by your ad" provision of BCRA, which requires "a statement by the candidate that identifies the candidate and states that the candidate has approved the communication."

("objective type"); or it may want the market to believe the company is underperforming to reap large profits ("biased type"). Biased A is concerned about his reputation, perhaps for legal reasons. Intermediary B is a financial analyst who issues reports about the company. Several recent lawsuits involve alleged uses of analysts to manipulate prices, which have led to Congressional investigations. For example, *Fortune* magazine reports: "Canadian insurer Fairfax Financial Holdings sues a group of hedge funds and research analysts for \$5 billion in New Jersey state court, alleging a stock market manipulation scheme in which the funds sold Fairfax's shares short, got analysts to write negative research reports that pushed the stock down, and made fortunes."[7]

II. Indirect Communication: Analysis

A. *Agenda pushing Equilibrium*

By assumption, objective A, B report truthfully: $m_A = s_A$, $m_B = m_A$. Biased A, B want to induce a high action to push their agenda and to appear objective. Biased B chooses a message to maximize his expected payoff given A's message:

$$\Pr(\eta = 1 \mid m_B) + \beta \mathbb{E}_\eta[\Pr(B = o \mid m_B, \eta) \mid m_A].$$

The first part, $\Pr(\eta = 1 \mid m_B)$, is C's optimal action given m_B, reflecting how far biased B can push his agenda. The second part reflects biased B's reputational concerns. In particular, $\Pr(B = o \mid m_B, \eta)$ is C's posterior belief of B's objectivity, given m_B and the observed state η. The expectation is taken with respect to state η conditional on m_A, because B knows only m_A when choosing m_B. Similarly, biased A chooses m_A to maximize his expected payoff:

$$\mathbb{E}m_B[\Pr(\eta = 1 \mid m_B) \mid m_A] + \alpha \mathbb{E}m_B[\mathbb{E}_\eta[\Pr(A = o \mid m_B, \eta) \mid s_A] \mid m_A].$$

Note that biased A takes expectation not only with respect to η given his signal s_A, but also with respect to m_B, which determines both C's action and A's posterior objectivity.

Before analyzing biased A, B's behavior, it helps to identify some key equilibrium properties.

DEFINITION 1: *In an "agenda pushing equilibrium," a biased agent reports his information truthfully if it supports his agenda, and with some probability if it does not: biased A reports* $m_A = 1$ *if* $s_A = 1$, *and* $m_A = 0$ *with probability* $x \in [0, 1]$ *if* $s_A = 0$; *biased B reports* $m_B = 1$ *if* $m_A = 1$, *and* $m_B = 0$ *with probability* $y \in [0, 1]$ *if* $m_A = 0$.

The corresponding strategies are called biased A, B's "agenda pushing strategy." An equilibrium is informative if C takes different actions based on B's message. Observe that, as long as the agents may be objective ($\theta_A > 0, \theta_B > 0$), an agenda pushing equilibrium is informative even if the biased agents lie completely ($x = y = 0$). The reason is that if $m_B = 0$, C knows $s_A = 0$; but if $m_B = 1$, C knows that $s_A = 1$ with some probability given the presence of objective agents. Since $m_B = 1$ is more indicative of $\eta = 1$ than $m_B = 0$, it leads C to take an action more favorable to biased agents' agenda.

[7] Bethany McLean, *Fortune* editor-at-large, "The inside story of a Wall Street battle royal," March 6, 2007.

LEMMA 1: *Every informative equilibrium is an agenda pushing equilibrium with* $x < 1$ *and* $y < 1$.

First, by ruling out $x = 1$ or $y = 1$, Lemma 1 shows that truth telling can not be supported in any informative equilibrium. In a putative truth telling equilibrium, message $m_B = 1$ is credible. Also, a biased agent's reputation is unaffected by his message because both types of agent report honestly. Thus the biased agent strictly gains from reporting in favor of his agenda, which is a contradiction. In contrast, $x = 0$ or $y = 0$ can be supported in an informative equilibrium because a message of 1 always induces a higher action than a message of 0, and thus a biased agent who cares little about reputation prefers reporting 1 regardless of his information. Further, if $x \in (0, 1)$ and $y \in (0, 1)$, then in an agenda pushing equilibrium, biased A is indifferent between reporting $m_A = 0$ and $m_A = 1$ if $s_A = 0$; and similarly, biased B is indifferent between $m_B = 0$ and $m_B = 1$ if $m_A = 0$. Intuitively, any benefit a biased agent gains from reporting one message over the other is exactly offset by the loss in his (expected) reputation.

Second, since the direction of bias is known, one may expect that a perverse equilibrium exists in which a biased agent intentionally distances himself from his agenda to appear objective. Lemma 1 shows that this is not the case. Suppose biased B understates evidence in favor of his agenda, then a message of 1 becomes *both* credible because it is more likely to come from an objective agent, and a positive sign of objectivity, which is a contradiction.

If the agents are unlikely to be objective ($\theta_A \leq 0.5$ and $\theta_B \leq 0.5$), there also exists an uninformative equilibrium in which C takes action $a = 0.5$ regardless of m_B.[8] This paper restricts attention to informative equilibria, which seem more sensible because both the biased agents and the decision maker receive lower payoffs in an uninformative equilibrium if it exists.

B. *Truth Telling Probabilities are Strategic Complements*

Given Lemma 1, the strategies of biased A and B can be represented by their respective truth telling probabilities x and y when the information does not support their agenda. To establish a key property of the best response of one biased agent to the other's truth telling probability, this subsection considers the net impact on a biased agent if he lies to push his agenda.

If $s_A = 0$, how is biased A affected if he reports $m_A = 1$ instead of $m_A = 0$? At first glance, this seems unclear because C does not observe A's message, whether it is associated with bias or not. After all, only B's message reaches C and influences her action directly. Let decision maker C's belief that agent A reports $m_A = 0$ given $s_A = 0$ be $N_x = \theta_A + (1 - \theta_A)x$; and similarly let her belief that agent B reports $m_B = 0$ given $m_A = 0$ be $N_y = \theta_B + (1 - \theta_B)y$. Interestingly, if $s_A = 0$, both biased A's agenda pushing benefit and his reputation cost from reporting $m_A = 1$ instead of $m_A = 0$ contain a common factor N_y. To see this, note that the pivotal event for biased A—which drives his message choice—is if C receives a different message from B *because of* A's message. Since biased B always reports $m_B = 1$ if $m_A = 1$, A's message makes a difference only if biased B truthfully reports $m_B = 0$ if $m_A = 0$. In particular, biased A's agenda pushing benefit if he reports $m_A = 1$ versus $m_A = 0$ is $N_y AP(x, y)$, where $AP(x, y)$ is the change in C's action if biased B reports $m_B = 1$ instead of $m_B = 0$:

$$(1) \qquad AP(x, y) \equiv \Pr(\eta = 1 \,|\, m_B = 1) - \Pr(\eta = 1 \,|\, m_B = 0) = \frac{p_A - 0.5}{1 - 0.5 N_x N_y}.$$

[8] In this equilibrium, m_A is uninformative because biased A randomizes in the opposite direction of s_A with sufficiently high probabilities to "overwhelm" the informative message from objective A. Thus agent B believes $s_A = 0$ and $s_A = 1$ are equally likely. Also, biased B randomizes in the opposite direction of m_A so that, at the evaluation stage, C believes $m_A = 0$ and $m_A = 1$ are equally likely, and thus A's reputation is unaffected by m_B.

This benefit is positive and increasing in x: the more truthful biased A is, the more likely C is swayed by B's message.

Biased A's reputation cost if he reports $m_A = 1$ instead of $m_A = 0$ depends on the observed true state at the evaluation stage. It is given by

$$\Pr(A = o \mid m_B = 0, \eta) - \Pr(A = o \mid m_B = 1, \eta),$$

which is always positive, regardless of the realized state. Given the agenda pushing strategies, if $m_B = 0$, C knows that $m_A = 0$, which is a positive sign of both A's and B's objectivity. If $m_B = 1$, then either agent may have distorted, and thus C downgrades both agents' posterior objectivity. At the information transmission stage, biased A's (expected) reputation cost given $s_A = 0$ is $N_y RC_A(x,y)$, where:

$$(2) \qquad RC_A(x,y) = \alpha\left[p_A \frac{N_y(1-N_x)\theta_A}{N_x(1-p_A N_x N_y)} + (1-p_A)\frac{N_y(1-N_x)\theta_A}{N_x(1-(1-p_A)N_x N_y)}\right].$$

The first term in the bracket corresponds to $\eta = 0$ and the second term corresponds to $\eta = 1$. Biased A's reputation cost decreases in x because the more truthful he is, the less C modifies her view of his objectivity from m_B.

Because C only observes m_B at the decision making stage, the agenda pushing benefit for biased B from reporting $m_B = 1$ instead of $m_B = 0$ is simply $AP(x,y)$ given in (1). Biased B's (expected) reputation cost, $RC_B(x,y)$, is symmetric to $RC_A(x,y)$. It is given by:

$$(3) \qquad RC_B(x,y) = \beta\left[p_A \frac{N_x(1-N_y)\theta_B}{N_y(1-p_A N_x N_y)} + (1-p_A)\frac{N_x(1-N_y)\theta_B}{N_y(1-(1-p_A)N_x N_y)}\right].$$

One implication is that biased A and B receive the same benefit AP from agenda pushing, *relative* to their respective reputation cost RC_A and RC_B. Any difference in their truth telling must be driven by differences in their reputation cost.[9]

For any y, the best response of biased A is defined as a truth telling probability x' such that: $x' = 0$ if $AP(x, y) \geq RC_A(x,y)$ for all x; $x' = 1$ if $AP(x, y) \leq RC_A(x,y)$ for all x; and otherwise any $x' \in (0,1)$ satisfying $AP(x', y) = RC_A(x', y)$.[10] Biased B's best response is analogously defined.

LEMMA 2: *There is a unique best response for each biased agent to any agenda pushing strategy of the other biased agent. Moreover, if positive, the best response is continuous and strictly increasing.*

The first part of Lemma 2 follows directly from the monotonicity of biased agents' agenda pushing benefit and their reputation cost with respect to their own truth telling probabilities. The second part shows that, perhaps surprisingly, biased B's truth telling has an unambiguous effect on biased A's truth telling: they are strategic complements. One may expect that the net effect

[9] This also implies that a biased agent's truth telling is independent of his location in a model with many intermediaries; see Proposition 6.

[10] If C believes that biased A reports truthfully with probability x', biased A is indifferent among any randomization between reporting $m_A = 1$ and $m_A = 0$. Thus any $x \in [0,1]$ is a best response here. However, only $x = x'$ is consistent with the definition of equilibrium, as $x \neq x'$ implies that biased A is no longer indifferent when C computes her action using x.

is ambiguous because, from (1) and (2), both AP and RC_A increase in y. The effect of changes in y on AP may be referred to as the *credibility reducing* effect of the intermediary, while the effect on RC_A may be referred to as the *blame sharing* effect. As y increases, m_B becomes more informative and thus agenda pushing by biased A becomes more effective. At the same time, C is more likely to attribute $m_B = 1$ to A's distortion than to B's, and thus it becomes more costly for biased A to lie. Although the net effect is generally ambiguous for any given x, y, using biased A's indifference condition between $m_A = 0$ and $m_A = 1$ if $x > 0$, the comparison between the credibility reducing effect and the blame sharing effect can be decomposed into two parts, corresponding to $\eta = 0$ and $\eta = 1$ at the evaluation stage. Since $s_A = 0$, biased A knows that $\eta = 0$ is more likely than $\eta = 1$, implying that the first part dominates.

In the first part of the comparison (conditional on $\eta = 0$), the credibility reducing effect is smaller than the blame sharing effect. For decision maker C, more truthful reporting from biased B (a higher y) matters only in the event that biased B distorts A's message: $m_B = 1$ but $m_A = 0$. Because if $m_B = 0$, C knows that $m_A = 0$ and $s_A = 0$, hence neither her action nor her evaluation of A is affected by a higher y. Also, biased B reports $m_B = 1$ with probability one if $m_A = 1$, in which case C is unaffected by any change in y. Note that C has different information at the decision making stage and at the evaluation stage. At the decision making stage, she knows only m_B, thus the impact of a higher y on AP is proportional to:

$$\Pr(m_A = 0 \mid m_B = 1) = \frac{0.5N_x(1 - N_y)}{1 - 0.5N_xN_y}.$$

But when she evaluates A given $\eta = 0$, the impact of a higher y on RC_A is proportional to:

$$\Pr(m_A = 0 \mid m_B = 1, \eta = 0) = \frac{p_A N_x(1 - N_y)}{1 - p_A N_x N_y}.$$

The first probability is smaller than the second, because without knowing $\eta = 0$, it is possible that $m_B = 1$ comes from $m_A = 0$ or $m_A = 1$; while given $\eta = 0$, it is more likely that $m_B = 1$ comes from $m_A = 0$ (and hence biased B's distortion). Consequently, as y increases, C revises her posterior belief about A's objectivity more than her action.

C. *The Unique Agenda Pushing Equilibrium*

An equilibrium of the indirect communication game is an intersection of biased A's and B's best responses. The following result establishes the existence and uniqueness of the equilibrium and provides a characterization. Existence follows from continuity of the best responses, and the uniqueness results from establishing that biased A's best response is always steeper than biased B's (if positive), which ensures that their best responses cannot cross more than once. Intuitively, a biased agent's reputation cost responds more to changes in his own truth telling probability than the other biased agent's reputation cost does.

PROPOSITION 1: *An agenda pushing equilibrium exists and is unique. In this equilibrium, there exist cutoff values α^i, β^i such that: (i) if a biased agent attaches a sufficiently high weight to his reputation, he reports truthfully with a positive probability: $x > 0$ if $\alpha > \alpha^i$ and $y > 0$ if $\beta > \beta^i$; (ii) if biased A, B both attach low weights to their reputations, they lie with probability one: $x = 0$ and $y = 0$ if $\alpha \leq \alpha^i$ and $\beta \leq \beta^i$.*

It is worth noting that if a biased agent has high reputational concerns, he reports information unfavorable to his agenda truthfully with some positive probability, but the reverse is not true. This is because α^i and β^i are defined such that $AP(0,0) = RC_A(0,0)$ at $\alpha = \alpha^i$ and $AP(0,0) = RC_B(0,0)$ at $\beta = \beta^i$.[11] If $\alpha > \alpha^i$, for example, biased A cannot afford to lie completely even if biased B does. Due to the complementarity of their best responses, biased A reports truthfully with a positive probability in equilibrium. If $\alpha \leq \alpha^i$, biased A can afford to lie completely if $y = 0$, but he may not if $y > 0$ due to their complementarity: either $x = 0$ or $x > 0$ is possible.

Changes in a biased agent's reputational concerns affect his truth telling. This is directly applicable to situations where one out of several agents may have leaked information to agent B, but A's exact identity is unknown. In this case, even if biased A has high reputational concerns, the effective α is smaller, and thus biased A is more apt to lie completely.[12] In the political campaign example, biased A, the candidate, may lie more if the voters know that A is only one of several possible sources of negative attacks against his opponent. Further, changes in α or β may affect all biased agents' truth telling probabilities because of their strategic complementarity. This matters to a decision maker who cannot reach all agents directly, perhaps because of existing anonymity granting rules of the media or laws protecting whistleblowers. For instance, several court decisions in recent years have grappled with setting appropriate legal guidelines for when intermediaries can be compelled to divulge the identities of their sources. In Doe versus Cahill (Del. 2005), the Delaware Supreme Court considered what a plaintiff must show in order to obtain a subpoena requiring an Internet service provider to disclose who posted anonymous comments online about a politician. In these legal examples, the financial cost of the intermediary (the Internet service provider) increases if he is more likely to be held liable for libel; while the cost for the source (bloggers) increases if the intermediary is more easily compelled to reveal his identity. The next result follows directly from Lemma 2 and Proposition 1.

COROLLARY 1: *Suppose that $\alpha > \alpha^i$ and $\beta > \beta^i$. Then biased A and B become more (less) truthful if either becomes more (less) concerned with his reputation: both x and y increase in α and in β.*

A biased media outlet reports more truthfully if it faces higher fines for granting anonymity too casually. Corollary 1 shows that this also makes it more costly for biased A to lie. Therefore the decision maker can improve the overall reporting accuracy by increasing the reputation cost of biased B. For example, the *New York Times* recently imposed a higher anonymity granting standard, because "the proliferation of critics and the growing public cynicism about the news media pose a threat to our authority and credibility that cannot go unanswered" (Bill Keller, executive editor of *New York Times*. "Assuring Our Credibility," memo, June 23, 2005). However, information also deteriorates quickly even if only one biased agent, say, a politician whose public life is drawing to an end, cares less about his reputation.[13]

[11] Although similar cutoffs can be defined for any fixed x and y, a sufficient condition for an interior equilibrium to exist is $\alpha > \alpha^i$ and $\beta > \beta^i$.

[12] Anonymity reduces biased A's reputation cost because he receives the same agenda pushing benefit, but pays his reputation cost only probabilistically, i.e., in the event C attributes the original message to him. In contrast, Kohei Kawamura (2006) shows that anonymity may improve a sender's truthful revelation. In his model, because the receiver gets multiple messages, his action is less responsive to each sender's message, which reduces a sender's benefit from lying.

[13] This result contrasts with Morris (2001), which shows that no information is communicated if both the objective and the biased advisor have sufficiently high reputational concerns. Because if the objective advisor avoids a message due to his higher reputational concerns, the biased advisor must avoid it as well. Here, the objective agents always report truthfully, thus only the biased agents report more truthfully when they face higher reputational concerns.

III. Indirect versus Direct Communication

A. *Direct Communication*

Suppose that A sends a message to C directly, which is equivalent to the case if biased B faces an infinitely high reputation cost ($\beta = \infty$), or if B is known to be objective ($\theta_B = 1$). As a special type of indirect communication, biased A's equilibrium behavior takes a special form of that in Section II.

PROPOSITION 2: *There exists a cutoff value α^d such that, in the unique equilibrium, biased A always reports $m_A = 1$ if $\alpha \leq \alpha^d$. If $\alpha > \alpha^d$, he reports $m_A = 0$ with probability $x^d > 0$ if $s_A = 0$.*

Biased A reports $m_A = 0$ truthfully sometimes if he cares sufficiently about his future reputation. Observe that the cutoff value of direct communication α^d, defined in the Appendix, is smaller than that of indirect communication α^i, because the presence of a possibly biased B makes it less costly for biased A to lie, everything else being equal. Also, in contrast to the indirect communication model, no uninformative equilibrium exists. Since C observes m_A and forms different beliefs about his objectivity at the evaluation stage, biased A must pay a reputation cost if he randomizes since he is more likely to be wrong than objective A. That is, by making m_A uninformative, biased A receives zero agenda pushing benefit, but pays a strictly positive reputation cost, which cannot be part of an equilibrium.

One may ask, in the earlier examples, whether a candidate lies less (against the opponent) if he is perceived to be very objective; or whether the government pushes its agenda less often if its private information becomes more accurate. The next result shows that biased A's truth telling is *nonmonotonic* in both his prior objectivity and the weight he attaches to his reputation.

COROLLARY 2: (*i*) *If $\alpha \leq p_A - 0.5$, biased A always reports $m_A = 1$. If $\alpha > p_A - 0.5$, then x^d first increases and then decreases in A's prior objectivity θ_A, becoming zero when θ_A is sufficiently high.* (*ii*) *If $\alpha \leq 1/(2 - \theta_A)$, then x^d first decreases in A's signal quality p_A and then becomes zero when p_A is sufficiently high. If α is sufficiently high, then $x^d > 0$, and it first decreases but eventually increases in p_A when p_A is sufficiently high.*

One may expect a politician with a good reputation at stake to be more truthful since he has more to lose. Instead, Corollary 2 shows that biased A is most truthful when his reputation is most responsive to his message, which occurs if his prior objectivity is in the intermediate range. If θ_A is sufficiently close to 0 and 1, A's message, right or wrong, has little impact on his reputation.[14]

More surprisingly, biased A may lie more, not less, as his information becomes more accurate. Suppose that α is sufficiently high such that $x^d > 0$ for all p_A. If p_A increases, biased A's agenda pushing has a stronger impact on C, which increases his incentive to lie; but biased A also shares more blame whenever $m_B = 1$ turns out wrong, which decreases his incentive to lie. Corollary 2 shows that if p_A is just above 0.5, even objective A is often wrong; and thus biased A's gain in agenda pushing dominates and he becomes less "fair and balanced" as his signal becomes more precise. When p_A becomes sufficiently high, however, a wrong message is (almost) a sure sign of bias, and thus the second effect dominates and biased A lies less.

[14] This is similar to Bénabou and Laroque (1992), who show that a biased agent has little incentive to invest in his reputation (report truthfully) when his existing reputation is sufficiently high or sufficiently low.

B. *Information Loss of Indirect Communication*

Sometimes agent A may communicate only in a particular way. In a government, officials may not be allowed to leak classified information to third parties: they must release information publicly, for example, through news conferences. Recall from propositions 1 and 2 that biased A reports $s_A = 0$ truthfully with probability x with intermediary B, and x^d without him. How does the channel of communication affect the decision maker's welfare?

COROLLARY 3: *Biased A lies less under direct communication than under indirect communication: $x^d \geq x$. The inequality is strict if $x^d > 0$.*

This result is a direct consequence of Corollary 1, which shows that both agents report less truthfully if biased B becomes less concerned with his reputation, as when biased A changes from direct to indirect communication. Adding an intermediary enables biased A to lie more, in addition to any distortions introduced by the intermediary.

The decision maker's expected payoff under indirect communication is:

$$\text{(4)} \qquad \mathbb{E}U_C = -\mathbb{E}_{m_B}[(\Pr(\eta = 1 \mid m_B) - \eta)^2] = -\frac{1}{2}\frac{[1 - (p_A^2 + (1 - p_A)^2)N_x N_y]}{2 - N_x N_y}.$$

This payoff strictly increases in x, y: the more truthful biased A and B are, the better off C is. It is immediate that C receives a lower expected payoff under indirect communication than direct communication because $x \leq x^d$, $y < 1$. A closer look at (4) shows that she receives a lower payoff under indirect communication for two reasons. First, because she receives a distorted message 1 with a higher probability: $\Pr(m_A = 1) < \Pr(m_B = 1)$. Second, she is also worse off because message 1 is less informative: $\Pr(\eta = 1 \mid m_A = 1) > \Pr(\eta = 1 \mid m_B = 1)$. In the opening example where the public chooses war or peace according to information learned through the media, not only the country is more likely to go to war on false grounds, it may also discount genuine threats. The decision maker thus prefers direct communication.

In a broader context, note that these communication channels induce different distributions of A's posterior objectivity. From an ex ante point of view, A's (expected) posterior objectivity is simply his prior θ_A in either channel by the law of iterated expectations. However, Corollary 3 suggests that indirect communication may provide more extreme estimates of A's type, which can be useful if learning A's objectivity is of value to C.

C. *Biased A's Ex Ante Choice of Communication Channels*

A government with an agenda to push faces a tradeoff: information transmitted directly is more credible while information transmitted through an intermediary is less costly in terms of his reputation. How should it balance these considerations and choose a communication channel given its own characteristics as well as those of the intermediary? This subsection studies biased A's ex ante channel choice, that is, before he observes signal s_A.[15]

[15] Ex ante choice is widely studied in the literature on information sharing among oligopolies, where information exchange decisions are taken prior to the arrival of private information, such as the realization of cost; see, for instance, Carl Shapiro (1986), David A. Malueg and Shunichi O. Tsutsui (1996), and the references within. Also, biased A's channel choice is one way for him to manipulate the informativeness of his signal, in particular C's belief about his objectivity. Leonard J. Mirman, Larry Samuelson, and Edward E. Schlee (1994) examine strategic manipulation of signal informativeness in a duopoly, where firms may adjust outputs away from myopically optimal levels to affect the informativeness of the market price.

To begin with, assume that objective A, who is nonstrategic, uses direct communication with probability $\mu \in (0, 1)$.[16] This assumption may be justified on institutional grounds; for example, many government agencies routinely give news conferences and background briefings to the media. It may also result from capacity constraints. For example, a political candidate can afford only a limited number of direct campaign advertisements. Note that from an ex ante perspective, which communication channel makes biased A better off is not obvious. In addition to the tradeoff between credibility and reputation, now the channel choice signals A's type in that the channel more likely chosen by biased A becomes a negative signal of A's objectivity, which affects how his message is interpreted and his reputation.

The channel choice game begins with a new stage in which agent A chooses a channel, either direct or indirect, which is observed by both B and C. The game then proceeds as described in Section I and thus all the results hold. The equilibrium is defined in the usual way in this new game. Let γ be the probability that biased A chooses direct communication. Then A's objectivity given his prior θ_A and the channel chosen can be computed by Bayes's rule. Denote A's interim objectivity given direct and indirect communication as θ_A^d and θ_A^i respectively.

Biased A's channel choice depends crucially on his tradeoff between message credibility and reputation. If α is sufficiently low, by propositions 1 and 2, biased A always reports $m_A = 1$ in either channel. Suppose that biased A uses direct communication with the same probability as objective A ($\gamma = \mu$), then no inference about his objectivity is made based on the channel chosen: $\theta_{\mathrm{A}}^d = \theta_A^i = \theta_A$. In this case, since B reduces the credibility of message 1, the expected action induced by $\mathrm{m}_A = 1$ is always higher than that induced by $m_B = 1$. Thus biased A wants to use direct communication more often than objective A: $\gamma > \mu$. The resulting negative inference about biased A's objectivity from doing so, due to his low reputational concerns, is outweighed by his gain in credibility from having no distortion introduced by B. Indeed, biased A may use direct communication exclusively: $\gamma = 1$. This occurs if B is perceived to be so biased that the loss in biased A's credibility is larger than the gain in his reputation even if he convinces B and C that he is objective by using indirect communication.

At the other extreme, if α is sufficiently high, biased A reports $m_A = 0$ with a positive probability in either channel by propositions 1 and 2. Then, by the law of iterated expectations, his ex ante expected payoff with direct and indirect communication is respectively $0.5 + \alpha\theta_A^d$ and $0.5 + \alpha\theta_A^i$, the sum of C's beliefs about the true state and about A's objectivity after observing the channel chosen, but before receiving the message.[17] As a result, biased A has to use direct communication with the same probability as objective A. If for instance $\gamma > \mu$, then $\theta_A^d < \theta_A^i$ and biased A would strictly prefer indirect communication, implying that $\gamma = 0$, which is a contradiction.

Biased A's equilibrium channel choice if α is in the intermediate range is more complicated to characterize for two reasons. First, his reporting strategy within a channel depends on θ_A^d and θ_A^i, which are endogenous in this model. Second, the strategic interactions between biased A and biased B also depend on the endogenous θ_A^i. To see a complication that arises from the latter, let $\mathbb{E}U_A^d(\theta_A^d)$ and $\mathbb{E}U_A^i(\theta_A^i)$ be biased A's expected payoff from direct and indirect communication given θ_A^d, θ_A^i. Note that $\mathbb{E}U_A^d(\theta_A^d)$ strictly increases in θ_A^d: the more objective biased A is perceived to be, the more effective is his agenda pushing and the better is his reputation.[18] However, $\mathbb{E}U_A^i(\theta_A^i)$ may not be monotonic in θ_A^i. Biased A's expected payoff increases in θ_A^i holding biased B's truth telling probability y fixed, but an increase in θ_A^i has an ambiguous effect on y: it increases both

[16] If $\mu = 0$ or 1, then biased A must choose the same channel that objective A does in equilibrium. Channel choice when objective A is also strategic is discussed further in Section VA.

[17] The law of iterated expectations applies with indirect communication because biased B's equilibrium behavior is already taken into account in biased A's indifference condition.

[18] This holds for any x^d even though x^d is nonmonotonic in θ_A^d (Corollary 2). By the Envelope Theorem, biased A responds optimally to changes in θ_A^d, and thus the change in x^d with respect to θ_A^d drops out in calculating $\mathbb{E}U_A^d(\theta_A^d)$.

biased B's agenda pushing benefit and his reputation cost as C now attributes more blame to B. Consequently, biased A's expected payoff may increase or decrease in θ_A^i.[19] The following result characterizes the unique equilibrium of the channel choice game in the special case when biased B always reports $m_B = 1$ ($y = 0$). Because biased B's truth telling probability does not interact with θ_A^i, $\mathbb{E}U_A^i(\theta_A^i)$ strictly increases in θ_A^i.

PROPOSITION 3: *If β is sufficiently low, then there exists a cutoff value $\alpha^s < \alpha^d$ such that in the unique equilibrium,* (1) *if $\alpha < \alpha^s$, biased A uses direct communication more often than objective A: $\mu < \gamma \leq 1$;* (2) *if $\alpha \in [\alpha^s, \alpha^i]$, biased A uses direct communication less often than objective A: $0 < \gamma \leq \mu$; and (iii) if $\alpha > \alpha^i$, then $\gamma = \mu$.*

Proposition 3 shows that, if the biased intermediary cares little about his reputation, biased A's equilibrium channel choice reverses itself in the intermediate range of α. Specifically, there exists a cutoff value α^s such that biased A chooses direct communication with a higher probability than objective A if $\alpha < \alpha^s$, but the reverse is true if $\alpha \in [\alpha^s, \alpha^i]$. To see why biased A may use indirect communication more often than objective A for some range of α, first recall biased A's equilibrium behavior in both communication channels when there is no channel choice. Since no inference is made about A's type, biased A is believed to be objective with probability θ_A in either channel. By propositions 1 and 2, if $\alpha < \alpha^d$, biased A always reports $m_A = 1$ in direct communication, and thus he also reports $m_A = 1$ via any intermediary: $x^d = 0, x = 0$. If $\alpha \in [\alpha^d, \alpha^i]$, then biased A reports $s_A = 0$ truthfully sometimes with direct communication, but he still reports $m_A = 1$ via an intermediary due to the blame sharing effect: $x^d > 0, x = 0$. In this case, biased A strictly prefers indirect communication to direct communication: $\mathbb{E}U_A^i(\theta_A) > 0.5 + \alpha\theta_A$ if $x = 0$ and $x^d > 0$. Given the monotonicity of both $\mathbb{E}U_A^d$ and $\mathbb{E}U_A^i$, biased A uses indirect communication with a greater probability than objective A in equilibrium when $\alpha \in [\alpha^d, \alpha^i]$. Clearly, biased A takes advantage of the fact that using an intermediary lowers his reputation cost to distort more than he does in direct communication. Since at $\alpha = \alpha^d$, biased A is strictly better off with indirect communication without channel choice, the critical value α^s at which biased A reverses his channel choice is lower than α^d.

A related question is how biased A's channel choice is affected by the characteristics of the intermediary he faces: a newspaper and a think tank may differ in their perceived objectivity.

COROLLARY 4: *If β is sufficiently low, and α is smaller than but sufficiently close to α^i, then biased A is less likely to use indirect communication if B is more objective: γ increases in θ_B.*

A more objective intermediary is more credible, but a poorer choice in terms of blame sharing. As α increases, however, the blame sharing effect of an intermediary becomes increasingly more important than the credibility reducing effect. Although generally ambiguous, the latter effect can outweigh the former if α is sufficiently close to α^i and β is sufficiently low. By using biased A's indifference condition between messages if $s_A = 0$, biased A's ex ante payoff is shown to decrease in θ_B. Since biased A's expected payoff from using direct communication is unchanged, he is more likely to use direct communication if the intermediary becomes more objective.

[19] For example, suppose that α is smaller than, but sufficiently close to, α^i, and β is sufficiently high such that $x = 0, y > 0$ in the relevant range of parameter values. Then $\mathbb{E}U_A^i$ decreases in y, which increases in θ_A^i, and thus biased A's expected payoff decreases if θ_A^i increases. Intuitively, because $x = 0$, the perceived probability that agent A reports truthfully $N_x = \theta_A^i$ increases with θ_A^i. Thus biased B's truth telling incentives also increase due to their strategic complementarity (y rises). Thus C attributes less blame to B's distortion than to A's. Because the blame sharing effect is stronger than the credibility reducing effect at α sufficiently close to α^i, biased A is worse off.

Corollary 4 suggests that agent B with a lower prior objectivity may be more likely to have access to biased A. Because biased B has no private information, he cannot influence C at all without A. If A is unlikely to be objective (θ_A sufficiently low) and β is sufficiently low, biased B's expected payoff is close to C's expected action:

$$0.5\gamma + \left(\frac{2p_A - 1}{2 - \theta_A^i \theta_B} + 1 - p_A\right)(1 - \gamma).$$

The first part is when A chooses direct communication and thus B is not used; while the second part is when B is used and he reports $m_B = 1$ due to his low reputational concerns. It is easy to see that biased B's expected payoff decreases in γ but increases in θ_B. That is, biased B prefers being used, which gives him a chance of pushing his agenda. But conditional on being used, the more objective B is perceived to be, the more credible his message is. If the first effect dominates, biased B prefers a lower θ_B since γ decreases in θ_B from Corollary 4. This provides a new rationale for media bias: B may prefer to appear more biased to encourage biased A to use him, and in turn becomes more influential.[20]

IV. Extensions

A. *Opposite Bias*

So far, the intermediary, if biased, prefers high actions just like the biased sender. In other possible applications, however, B may be biased in different ways. In the opening example, a prowar government may face a media outlet with a possible antiwar bias; in military settings, the sender may have to communicate through unfriendly intermediaries.[21] A question of practical importance, and the first departure from the main model, is whether biased A distorts more or less to push his agenda if B may have an opposite bias. Assume that biased B prefers low actions regardless of the state, and his payoff function is:

$$U_B = -a + \beta\pi_B.$$

All other assumptions remain. An agenda pushing equilibrium is defined such that biased A behaves in the same way as in Section II, while biased B reports $m_B = 0$ if $m_A = 0$; and $m_B = 1$ with probability $w \in [0, 1]$ if $m_A = 1$. The corresponding strategies are biased A and B's agenda pushing strategies.

Given the agenda pushing strategies, biased B may report $m_B = 0$ when $m_A = 1$, thus $m_B = 1$ becomes a sure sign of $m_A = 1$. Consequently, $m_B = 1$ is a better sign of agent B's objectivity than $m_B = 0$. This contrasts with the same-biased model in which $m_B = 0$, as a sure sign of $m_A = 0$, is a better sign of B's objectivity. Despite this difference, there are similarities with the same-biased model. First, $m_B = 1$ is still more indicative of $\eta = 1$ than $\eta = 0$. Second, $m_B = 0$ remains a better sign of A's objectivity than $m_B = 1$, because $m_B = 0$ results from $m_A = 0$ with a positive probability while $m_B = 1$ must come from $m_A = 1$. As in Lemma 1, any informative equilibrium in the opposite-biased model is an agenda pushing equilibrium in which biased agents lie to some extent to push their respective agenda.

[20] Terence Lim (2001) documented this effect in corporate earnings forecasting. In his model, analysts are assumed to learn more about a company from the management if they publish reports with a bias favorable to the company.

[21] For instance, according to some military officials, "the Pentagon is developing plans to provide news items, possibly even false ones, to foreign media organizations as part of a new effort to influence public sentiment and policy makers in both friendly and unfriendly countries." James Dao and Eric Schmitt, "Pentagon Readies Efforts to Sway Sentiment Abroad," *New York Times*, February 19, 2002.

PROPOSITION 4: *An informative equilibrium is an agenda pushing equilibrium, and exists. In any such equilibrium,* (1) *if* α *and* β *are sufficiently low, biased A and B lie with probability one:* $x = 0$ *and* $w = 0$*;* (2) *if* α *and* β *are sufficiently high, both biased A and B report truthfully with positive probabilities:* $x, w \in (0, 1)$*;* (3) *if positive, x and w are strategic substitutes.*

In the opposite-biased model, biased A and B's truth telling probabilities, if strictly positive, are strategic substitutes. If an antiwar media outlet reports the government's message more truthfully (w increases), a prowar government distorts more to push its agenda (x decreases), and vice versa. Although this result is the exact opposite of the strategic complements result in the main model, an increase in biased B's truth telling probability w has the same two effects on biased A as an increase in y. First, the credibility reducing effect remains because agent B has no private information but a possible bias. Observe that the agenda pushing benefit for biased A, $\Pr(\eta = 1 \mid m_B = 1) - \Pr(\eta = 1 \mid m_B = 0)$, is equal to biased B's benefit $\Pr(\eta = 0 \mid m_B = 0) - \Pr(\eta = 0 \mid m_B = 1)$. It increases in w because the more truthful biased B is, the more credible $m_B = 0$ is. Second, the blame sharing effect remains because C always attributes some blame to B if m_B is associated with B's possible bias. If biased B is more truthful, $m_B = 0$ becomes a better signal of A's objectivity while biased A's reputation given $m_B = 1$ is unaffected. Thus biased A's reputation cost increases in w as in the same-biased model.

Biased agents' truth telling probabilities become strategic substitutes in the opposite-biased model due to a change in the sign of the net effect, because the critical event changes from biased B distorting $m_A = 0$, to his distorting $m_A = 1$. The analysis follows the same steps as in the main model. First, using biased A's indifference condition between $m_A = 0$ and $m_A = 1$ when $s_A = 0$, the net effect of an increase in w on biased A can be decomposed into two parts, corresponding to $\eta = 0$ and $\eta = 1$ at the evaluation stage. Second, the first part dominates because biased A knows that $\eta = 0$ is more likely than $\eta = 1$ since $s_A = 0$. Third, the critical event is $m_A = 1$ but $m_B = 0$, because from the decision maker's point of view, more truthful reporting from biased B (a higher w) matters only in the event biased B distorted. Thus, the credibility reducing effect of an increase in w is proportional to $\Pr(m_A = 1 \mid m_B = 0)$, and the blame sharing effect is proportional to $\Pr(m_A = 1 \mid m_B = 0, \eta = 0)$. The second probability is smaller than the first, because if the true state is 0, it is less likely that $m_A = 1$. Since C changes her posterior belief about A's objectivity less than changing her action, biased A faces a lower reputation cost and reports less truthfully: x decreases in w.[22]

Propositions 1 and 4 have interesting implications in terms of biased A's behavior. Consider the case where β is so high that biased B is very truthful regardless of the direction of his bias (y or w sufficiently close to 1). Then biased A reports more truthfully if the intermediary may have an opposite bias than the same bias. Both reduce the credibility of A's message, but the same-biased intermediary is more effective in sharing biased A's blame due to their strategic complementarity. One hastens to add that C is not necessarily better off in the opposite-biased case, in which both signals may be distorted, while in the same-biased model only $s_A = 0$ is.[23]

[22] Maxim Ivanov (forthcoming) considers a model with a biased intermediary (mediator) in the Crawford and Sobel (1982)framework and shows that using an intermediary with a bias opposite to that of the sender can improve communication outcome. This is because the biased mediator adds endogenous noise by adjusting the message in the opposite direction of the sender's bias, similar to the agenda pushing strategy here. The focus of Ivanov (forthcoming), as well as Krishna and Morgan (2004); Blume, Board, and Kawamura (2007); and Goltsman et al. (2009), is on how to introduce noise to improve communication efficiency. This model focuses instead on the tradeoff between credibility and reputation cost when agents have reputational concerns, and shows that the decision maker may be worse off due to the blame sharing effect. See also the next subsection for a comparison with the noisy communication model.

[23] For instance, C's ex ante payoff in the opposite-biased model may be lower if biased A is very truthful due to his high reputational concerns, in which case the possible distortion of $m_A = 1$ leads C to take too high an action after $m_B = 0$.

B. *Noisy Intermediary*

Possible distortions from a strategic intermediary add endogenous noise to the message the decision maker receives. To isolate the effect of noise on biased A, the second departure from the main model is to assume that agent B is unbiased but "careless." Trying to pass on m_A truthfully, he may make a mistake and report the opposite with a small probability. Formally, for each $m_A \in \{0, 1\}$, biased B's message $m_B = m_A$ with probability y_n. All other assumptions remain.

Noise has two (by now) familiar effects on biased A. Biased A's agenda pushing benefit increases in y_n because the smaller the noise is, the more credible m_B is, increasing biased A's incentive to lie. Also, biased A's (expected) reputation cost increases in y_n due to the blame sharing effect, because the smaller the noise is, $m_B = 0$ is a better sign of A's objectivity and $m_B = 1$ is a worse sign, decreasing his incentive to lie. The net effect determines whether he lies more or less as the noise diminishes.

PROPOSITION 5: *If the noise is sufficiently small (y_n is sufficiently close to 1), an agenda pushing equilibrium exists in which biased A reports $m_A = 0$ if $s_A = 0$ with probability $x_n \in [0, 1)$. If positive, x_n increases in y_n if A's prior objectivity θ_A or if α is sufficiently high; and x_n decreases in y_n otherwise.*

Because the noise affects biased A when $m_A = 0$ and when $m_A = 1$ independently and symmetrically, one can separate its net effect into two parts. Holding biased A's behavior fixed, the net effect of an increase in y_n on biased A if $m_A = 0$ is the same as that with a same-biased intermediary who reports $m_B = 0$ if $m_A = 0$ with probability $y = y_n$; and the net effect on biased A if $m_A = 1$ is the same as that with an opposite-biased intermediary who reports $m_B = 1$ if $m_A = 1$ with probability $w = y_n$. More precisely, given x_n, y_n, if the noise diminishes (y_n increases), the net effect of a nonstrategic intermediary on biased A is the sum of the negative net effect of a same-biased intermediary (as y increases); and the positive net effect of an opposite-biased intermediary (as w increases). One way to think about the role of noise is that it works against the strategic effect by dampening the strategic interactions between biased agents. For instance, viewing through the lens of a same-biased intermediary, if biased B reports more truthfully, he shares less blame. However, the noise weakens this effect, and thus biased A's truth telling probability increases less (or even decrease) than without noise.

If biased A has a high prior objectivity or if his reputational concerns are sufficiently high ($\theta_A + (1 - \theta_A)x_n > 2 - \sqrt{2}$), biased A distorts more if there is more noise (x_n decreases as y_n decreases). Because C is more likely to attribute any wrong message to noise than to biased A's distortion, without the result is qualitatively the same as in the same-biased model. If instead, agent A is very biased and his reputational concerns are low, his message is heavily distorted even with no noise. The noise further reduces the message's credibility without significantly lowering biased A's reputation cost, and thus he needs to report more truthfully as noise increases (x_n increases as y_n decreases). The result is qualitatively the same as in the opposite-biased model.

The result that biased A may lie more or less with a noisy intermediary differs from recent papers such as Krishna and Morgan (2004) and Blume, Board, and Kawamura (2007), who show that adding a small amount of noise may improve biased A's truth telling incentives. In this model, noise still reduces the credibility of A's message, and thus biased A has an incentive to report more truthfully. Unlike the aforementioned papers, however, noise also reduces biased A's reputation cost which may lead to greater distortions if *a priori*, the sender's message is highly credible.

C. *Multiple Intermediaries*

In the third departure from the main model, this subsection considers the case of multiple intermediaries. Suppose that there are k agents, each of whom sends a message to his immediate successor, and decision maker $k + 1$ takes an action based on the last message m_k. For simplicity, let the agents be symmetric: $\theta_i = \theta$ and $\alpha_i = \alpha$ for all $i = 1, 2, \ldots, k$. Agent 1, the first agent, receives a private signal s_1: $\Pr(s_1 = \eta) = p_1 > 0.5$; all other agents are uninformed. The agents, if biased, all favor action $a = 1$. All other assumptions remain.

PROPOSITION 6: *There exists an agenda pushing equilibrium in which all biased agents report $m_i = 1$ if α is sufficiently low; otherwise, biased i reports $m_i = 0$ with the same probability $x_k \in (0, 1)$. Moreover, each biased agent lies with a higher probability as the number of agents increases: x_k decreases in k.*

As in the main model, the pivotal event for biased i is if $m_{i-1} = 0$, he can change the final message from $m_k = 0$ to $m_k = 1$ by distortion. Also, each biased i's agenda pushing benefit is $\Pr(\eta = 1 \mid m_k = 1) - \Pr(\eta = 1 \mid m_k = 0)$, multiplied by the factor $\Pr(m_k = 0 \mid m_i = 0)$; and his reputation cost is multiplied by the same factor. Therefore, biased i's truth telling depends solely on his reputation cost, which is a function of θ_i and α_i. Since these parameters are identical by assumption, all biased agents report truthfully with the same probability, regardless of their locations.[24] Moreover, these truth telling probabilities are strategic complements, implying that an additional agent reduces each biased agent's truth telling: $x_k = 0$ if k is sufficiently large.

V. Discussion and Concluding Remarks

A. *The Objective Agent Assumption*

In this model, an objective agent is assumed to report honestly. But even an objective agent who cares about the decision maker's welfare may also face reputational concerns. Formally, suppose that objective agents have the following payoff functions, where α^o, β^o are the respective weights objective A, B attach to their reputation:

$$U_A^o = -(a - \eta)^2 + \alpha^o \pi_A; \quad \text{and} \quad U_B^o = -(a - \eta)^2 + \beta^o \pi_B.$$

If $\alpha^o = 0$, $\beta^o = 0$, this model is similar to Sobel (1985) in that objective agents have exactly the same preferences as the decision maker. In this case, given a communication channel, the agenda pushing equilibrium found in this paper remains valid. In equilibrium, a message is always more likely to be correct than not: $\Pr(\eta = 1 \mid m_A = 1) > \Pr(\eta = 1 \mid m_A = 0)$ for the intermediary; and $\Pr(\eta = 1 \mid m_B = 1) > \Pr(\eta = 1 \mid m_B = 0)$ for the decision maker. Further, given this equilibrium behavior of B and C, objective A reports s_A truthfully; similarly, objective B reports truthfully as in the main model.

If $\alpha^o > 0$, $\beta^o > 0$, this model is similar to Morris (2001) in that the objective agents also care about reputation. In the opening example, such an objective government wants to inform

[24] A similar pivotal argument has been used in Eddie Dekel and Michele Piccione (2000) and Hao Li, Sherwin Rosen, and Wing Suen (2001) to show that the order of voting does not affect the voting outcomes in equilibrium. In a nonstrategic social network context, Peter DeMarzo, Dimitri Vayanos, and Jeffrey Zwiebel (2003) show that one's influence on other people depends not only on his information, but also his position in a given social network, because the agents do not account for possible repetitions in the information that reaches them.

the public about whether a military threat exists; but it does not want to look like a biased, warmongering hawk. For a given communication channel, the argument above shows that if α^o and β^o are sufficiently low, the equilibrium remains because the objective agents report truthfully by continuity. But if objective A or B has sufficiently high reputational concerns, any equilibrium is uninformative in a given communication channel. Suppose that β^o is sufficiently high; then as long as m_A is informative, there is a message that biased B would use to push his agenda. Objective B would avoid this message out of reputational concerns, and thus biased B cannot use this message either. Since both types of B use the same message, the equilibrium is uninformative. This is similar to Proposition 2 of Morris (2001), who shows in a model without intermediary that no informative equilibrium exists if the objective type has high reputational concerns.

A more subtle question is how such an objective A chooses a communication channel prior to receiving any information. The analysis so far suggests that, on one hand, direct communication such as a news conference is the preferred choice of a biased, prowar government who cares little about reputation, and thus an objective government may want to resort to indirect communication to avoid looking biased. On the other hand, any information leaked to the media may be distorted, misleading the public into more extreme actions. Web Appendix B shows that in a model of objective but strategic A, first, there always exist two pooling equilibria in which both types of A choose the same communication channel.[25] If objective A uses one channel exclusively, then biased A loses all reputation and his message is not credible if he uses the other channel.[26]

Web Appendix B considers whether other informative equilibria exist in which both channels are used. The answer turns out to be "no" if objective A has *no* reputational concerns. For such an objective A to use both channels, he must be ex ante indifferent between the actions induced by these channels. This implies that C must receive message 1 with the same probability in equilibrium in either channel, and take the same action after message 1. If biased A uses one channel exclusively in the putative equilibrium, the expected actions after message 1 differ and thus objective A cannot be indifferent. If instead, objective A is indifferent between the two channels, then biased A's reputational concerns alone determine his channel choice, which favor indirect communication because of the blame sharing effect. But if biased A uses indirect communication exclusively, then objective A strictly prefers direct communication. Thus no informative equilibria exist in which both communication channels are used. If objective A has low reputational concerns, however, informative equilibria in which both channels are used may exist. Web Appendix B shows one such example in which objective A uses both channels while biased A uses direct communication only. In this equilibrium, biased A cares little about reputation and is willing to use direct communication for the higher action he can induce, while objective A uses both channels because indirect communication gives him a higher reputation than direct. Intuitively, the objective agent's reputational concerns imply that his channel choice is no longer solely based on the action his message induces.

B. *Information Aggregation and Informed Intermediaries*

Agent B is uninformed in the current model, but he may have good information of his own in many marketing and medicine settings. Wei Li (forthcoming) considers the case where an objective agent reports the best information available, and studies how a source (such as a pharmaceutical company) tries to use well informed experts (such as physicians) to influence the decision maker (such as the patients). Although using one's own information is a sign of objectivity, the

[25] These pooling equilibria are supported by the out-of-equilibrium-path belief that A is biased if he deviates to the other channel. Also, both equilibria survive belief refinement such as D1 criterion for some parameter values.

[26] Similarly, no equilibrium exists in which objective A uses one channel exclusively and biased A uses both.

biased intermediary selectively incorporates the sender's information to push his agenda. The intermediary's truth telling always decreases in the sender's. Hence measures raising the sender's reputation cost, and thus his truth telling, may make the decision maker strictly worse off.

C. *Conclusion*

In the opening example, a government with an agenda to push may communicate with the public directly or through the media. The media shares the government's blame of releasing inaccurate information, thereby reducing the government's reputation cost. But indirect messages are less credible. The government is more likely to inform the public directly if it cares little about its reputation; but it is more likely to leak to the intermediary if it has moderate reputational concerns. The media outlet may benefit from appearing biased, because he is more likely to receive information from an agenda pushing government.

Little is understood about why different communication protocols coexist in organizations. Existing papers focus primarily on how firms aggregate and process information through direct communication (Kenneth J. Arrow and Roy Radner 1979; Kenneth J. Arrow 1985). This paper, however, suggests that indirect communication serves a useful purpose, and thus may be chosen to convey certain types of information. Also, indirect communication matters in studying misinformation in military operations. Direct communication is unlikely to be effective in this context, because many such operations are zero-sum by nature (Vincent P. Crawford 2003; Ken Hendricks and R. Preston McAfee 2006).

Appendix

PROOF OF LEMMA 1:

Let biased A report $m_A = 0$ with probability x if $s_A = 0$; and $m_A = 1$ with probability z if $s_A = 1$. Also, let biased B report $m_B = 0$ with probability y if $m_A = 0$; and $m_B = 1$ with probability w if $m_A = 1$. Denote the decision maker's beliefs that agent A and B report truthfully as: $N_x \equiv \theta_A + (1 - \theta_A)x$, $N_z \equiv \theta_A + (1 - \theta_A)z$, $N_y \equiv \theta_B + (1 - \theta_B)y$, $N_w \equiv \theta_B + (1 - \theta_B)w$.

Given the above strategies, consider biased B's truth telling incentives first. Let $\nu_1 \equiv \Pr(\eta = 1 \mid m_A = 1)$, $\nu_0 \equiv \Pr(\eta = 1 \mid m_A = 0)$, then:

$$\nu_1 = \frac{p_A N_z + (1 - p_A)(1 - N_x)}{1 - N_x + N_z}; \quad \text{and} \quad \nu_0 = \frac{p_A(1 - N_z) + (1 - p_A)N_x}{1 - N_z + N_x}.$$

First, decision maker C's actions given B's messages are respectively:

$$a_1^B \equiv \Pr(\eta = 1 \mid m_B = 1) = \frac{[p_A N_z + (1 - p_A)(1 - N_x)]N_w + [p_A(1 - N_z) + (1 - p_A)N_x](1 - N_y)}{(N_z + 1 - N_x)N_w + (1 - N_z + N_x)(1 - N_y)};$$

$$a_0^B \equiv \Pr(\eta = 1 \mid m_B = 0) = \frac{[p_A N_z + (1 - p_A)(1 - N_x)](1 - N_w) + [p_A(1 - N_z) + (1 - p_A)N_x]N_y}{(N_z + 1 - N_x)(1 - N_w) + (1 - N_z + N_x)N_y}.$$

Moreover, $a_1^B - a_0^B$ is proportional to, and has the same sign as, $(\nu_1 - \nu_0)[N_w N_y - (1 - N_w) \times (1 - N_y)]$.

Next, let $\pi_{B(m_B,\eta)} \equiv \Pr(B = o \mid m_B, \eta)$ be B's posterior objectivity conditional on m_B, η. Similarly, let $\pi_{B(m_B)} \equiv \Pr(B = o \mid m_B)$. Then we have:

$$\pi_{B(1,0)} = \frac{[(1-p_A)N_z + p_A(1-N_x)]\theta_B}{[(1-p_A)N_z + p_A(1-N_x)]N_w + [(1-p_A)(1-N_z) + p_AN_x](1-N_y)};$$

$$\pi_{B(0,0)} = \frac{[(1-p_A)(1-N_z) + p_AN_x]\theta_B}{[(1-p_A)(1-N_z) + p_AN_x]N_y + [(1-p_A)N_z + p_A(1-N_x)](1-N_w)};$$

$$\pi_{B(0,1)} = \frac{[p_A(1-N_z) + (1-p_A)N_x]\theta_B}{[p_A(1-N_z) + (1-p_A)N_x]N_y + [p_AN_z + (1-p_A)(1-N_x)](1-N_w)};$$

$$\pi_{B(1,1)} = \frac{[p_AN_z + (1-p_A)(1-N_x)]\theta_B}{[p_AN_z + (1-p_A)(1-N_x)]N_w + [p_A(1-N_z) + (1-p_A)N_x](1-N_y)}.$$

For biased B to report $m_A = 0$ and $m_A = 1$ truthfully, the following two incentive constraints (IC) must hold:

$$\text{(A1)} \qquad a_1^B - a_0^B \le \beta[(1-\nu_0)\pi_{B(0,0)} + \nu_0\pi_{B(0,1)} - (1-\nu_0)\pi_{B(1,0)} - \nu_0\pi_{B(1,1)}];$$

$$\text{(A2)} \qquad a_1^B - a_0^B \ge \beta[(1-\nu_1)\pi_{B(0,0)} + \nu_1\pi_{B(0,1)} - (1-\nu_1)\pi_{B(1,0)} - \nu_1\pi_{B(1,1)}].$$

Note that the LHS of IC (A1) and (A2) are the same, and the difference in their RHS is $\beta(\nu_1 - \nu_0)[\pi_{B(0,0)} + \pi_{B(1,1)} - \pi_{B(1,0)} - \pi_{B(0,1)}]$. Moreover, $\pi_{B(0,0)} > (\le)\pi_{B(0,1)}$ and $\pi_{B(1,1)} > (\le)\pi_{B(1,0)}$ if $\nu_1 > (\le)\nu_0$, and the equalities hold if $\nu_1 = \nu_0$. Together, this implies that the RHS of IC (A1) is (weakly) larger than the RHS of (A2): biased B pays a higher reputation cost if his message differs from the state more likely to be true given m_A. Intuitively, $\pi_{B(0,0)}, \pi_{B(1,1)}$ are B's posterior objectivity if his message is accurate, and $\pi_{B(1,0)}, \pi_{B(0,1)}$ are those if his message is inaccurate. If $\nu_1 > \nu_0$, then objective B is more accurate since he follows m_A; but if $\nu_1 < \nu_0$, then objective B is less accurate, in which case inaccurate messages lead to higher posterior objectivity.

Let $\pi_{A(m_B,\eta)} \equiv \Pr(A = o \mid m_B, \eta)$ be agent A's posterior objectivity given m_B, η. Similarly, let $\pi_{A(m_B)} \equiv \Pr(A = o \mid m_B)$, then:

$$\pi_{A(0,1)} = \frac{[p_A(1-N_w) + (1-p_A)N_y]\theta_A}{p_A[N_z(1-N_w) + (1-N_z)N_y] + (1-p_A)[N_xN_y + (1-N_x)(1-N_w)]};$$

$$\pi_{A(0,0)} = \frac{[p_AN_y + (1-p_A)(1-N_w)]\theta_A}{p_A[N_xN_y + (1-N_x)(1-N_y)] + (1-p_A)[N_z(1-N_w) + (1-N_z)N_y)]};$$

$$\pi_{A(1,0)} = \frac{[p_A(1-N_y) + (1-p_A)N_w]\theta_A}{p_A[N_x(1-N_y) + (1-N_x)N_w] + (1-p_A)[N_zN_w + (1-N_z)(1-N_y)]};$$

$$\pi_{A(1,1)} = \frac{[p_AN_w + (1-p_A)(1-N_y)]\theta_A}{p_A[N_zN_w + (1-N_z)(1-N_y)] + (1-p_A)[N_x(1-N_y) + (1-N_x)N_w]}.$$

If $s_A = 0$, the difference in biased A's expected payoffs between $m_A = 1$ and $m_A = 0$ is:

$$\mathbb{E}U_A(m_A = 1 \mid s_A = 0) - \mathbb{E}U_A(m_A = 0 \mid s_A = 0)$$

$$= N_w[a_1^B + \alpha p_A \pi_{A(1,0)} + \alpha(1-p_A)\pi_{A(1,1)}] + (1-N_w)[a_0^B + \alpha p_A \pi_{A(0,0)} + \alpha(1-p_A)\pi_{A(0,1)}]$$

$$- (1-N_y)[a_1^B + \alpha p_A \pi_{A(1,0)} + \alpha(1-p_A)\pi_{A(1,1)}] - N_y[a_0^B + \alpha p_A \pi_{A(0,0)} + \alpha(1-p_A)\pi_{A(0,1)}]$$

$$= (N_w - (1-N_y))[a_1^B - a_0^B + \alpha p_A(\pi_{A(1,0)} - \pi_{A(0,0)}) + \alpha(1-p_A)(\pi_{A(1,1)} - \pi_{A(0,1)})].$$

Clearly, both biased A's agenda pushing benefit and his reputation cost are multiplied by a common factor: $N_w - (1 - N_y)$. If $N_w - (1 - N_y) > 0$, then we can take out this factor, and biased A derives the same agenda pushing benefit, $a_1^B - a_0^B$, relative to his reputation cost as biased B. Biased A also faces two truth telling ICs given $s_A = 0$ and $s_A = 1$ respectively:

$$\text{(A3)} \qquad a_1^B - a_0^B \leq \alpha[p_A\pi_{A(0,0)} + (1-p_A)\pi_{A(0,1)} - p_A\pi_{A(1,0)} - (1-p_A)\pi_{A(1,1)}];$$

$$\text{(A4)} \qquad a_1^B - a_0^B \geq \alpha[(1-p_A)\pi_{A(0,0)} + p_A\pi_{A(0,1)} - (1-p_A)\pi_{A(1,0)} - p_A\pi_{A(1,1)}].$$

The difference in the RHS of (A3) and (A4) is $\alpha(2p_A - 1)[\pi_{A(0,0)} + \pi_{A(1,1)} - \pi_{A(0,1)} - \pi_{A(1,0)}]$.

We now consider the case when $\nu_1 \neq \nu_0$. Truth telling is impossible: if $y = w = 1$, then the LHS of IC (A1) and (A2) is nonzero, but the RHS of both are zero, and thus one of the ICs cannot hold. Recall from the analysis of IC (A1) and IC (A2) above, the LHS of (A1) and (A2) are the same, while the RHS of (A1) is larger than that of (A2). Hence if IC (A1) does not hold or hold with equality, IC (A2) must hold strictly. This rules out equilibrium strategy $y = 0$, $w \in [0, 1)$ and $w = 0$, $y \in [0, 1)$. It also rules out $y \in (0, 1)$, $w \in (0, 1)$ because biased B's ICs cannot both hold with equality. If m_A is informative, then only two strategies are possible: $y = 1$, $w \in [0, 1)$ and $y \in [0, 1)$, $w = 1$. In both cases, $N_w N_y - (1 - N_w)(1 - N_y) > 0$, hence $a_1^B - a_0^B$ has the same sign as $(\nu_1 - \nu_0)$. Also, since either $N_y = 1$ or $N_w = 1$, $N_w - (1 - N_y) > 0$, m_B affects biased A's expected payoff.

CLAIM 1: *If $\nu_1 \neq \nu_0$, biased B uses his agenda pushing strategy in equilibrium:* $y \in [0, 1)$, $w = 1$.

Suppose that $\nu_1 > \nu_0$, then $a_1^B - a_0^B > 0$. If in equilibrium, y = 1, w ∈ [0, 1), then we can show that $\pi_{B(1,0)} = \pi_{B(1,1)}$; and $\pi_{B(1,0)} > \pi_{B(0,0)}$, $\pi_{B(1,1)} > \pi_{B(0,1)}$. Thus biased B's reputation cost, the RHS of IC (A1), is negative while the LHS is positive, which violates IC (A1). Thus if $\nu_1 > \nu_0$, which is true if biased A uses any strategy involving x = 1 or z = 1, the only possible equilibrium is informative and it involves agenda pushing for B.

If $\nu_1 < \nu_0$, then $N_x N_z - (1 - N_x)(1 - N_z) < 0$, which requires biased A to distort both signals with positive probabilities. If in equilibrium, $y = 1$, $w \in [0, 1)$, then we can show that for biased A, $\pi_{A(1,1)} > \pi_{A(1,0)}$ and $\pi_{A(0,0)} > \pi_{A(0,1)}$. This implies that the RHS of IC (A3) is strictly larger than the RHS of (A4). Clearly, mixed strategies $x \in (0, 1)$, $z \in (0, 1)$ are impossible. Also, if IC (A3) does not hold or if it holds with equality, IC (A4) must hold strictly, which rules out $x \in [0, 1)$, $z = 0$. The only possible strategies for biased A are $x = 1$, $z \in [0, 1)$ and $x \in [0, 1)$, $z = 1$. But in both these cases, $\nu_1 > \nu_0$, a contradiction. Thus if $\nu_1 \neq \nu_0$, the only possible equilibrium strategy for biased B is the agenda pushing one: $y \in [0, 1)$, $w = 1$.

CLAIM 2: *If biased B uses the agenda pushing strategy, biased A must use agenda pushing strategy in equilibrium:* $x \in [0,1)$, $z = 1$.

If $y \in [0,1)$, $w = 1$, we can show that for biased A, $\pi_{A(1,1)} > \pi_{A(1,0)}$ and $\pi_{A(0,0)} > \pi_{A(0,1)}$, which rules out all biased A's strategies except for $x = 1$, $z \in [0,1)$ and $x \in [0,1)$, $z = 1$. If $x = 1$, $z \in [0,1)$, however, then the LHS of IC (A3) is positive, but the RHS is negative because $\pi_{A(0,1)} < \pi_{A(1,1)}$, $\pi_{A(0,0)} < \pi_{A(1,0)}$. Thus IC (A3) cannot hold, which is a contradiction. Hence biased A would deviate and report $m_A = 1$. This shows that if $\nu_1 \neq \nu_0$, the only possible equilibrium is: $x \in [0,1)$, $z = 1$ and $y \in [0,1)$, $w = 1$.

LEMMA A1: *If* $\theta_A \leq 0.5$ *and* $\theta_B \leq 0.5$, *there exists an uninformative equilibrium in which both biased A and B randomize so that B's message is useless:* $a_1^B = a_0^B = 0.5$.

PROOF:

See Web Appendix A.

PROOF OF LEMMA 2:

Given Lemma 1, we limit our attention to agenda pushing strategies. Recall that biased B has two truth telling ICs: (A1) and (A2). Also, the only IC that may hold with equality in equilibrium is IC (A1). Similarly, biased A has two truth telling ICs: (A3) and (A4), and the only IC that may hold with equality is IC (A3). To simplify notations, define the following functions of x, y:

$$\text{(A5)}\quad \xi(x,y) \equiv \frac{2p_A - 1}{2 - N_xN_y} - \alpha\theta_A\left[\frac{1}{N_x} - \frac{p_A(1-p_AN_y)}{1-p_AN_xN_y} - \frac{(1-p_A)[1-(1-p_A)N_y]}{1-(1-p_A)N_xN_y}\right];$$

$$\text{(A6)}\quad \psi(x,y) \equiv \frac{2p_A - 1}{2 - N_xN_y} - \beta\theta_B\left[\frac{1}{N_y} - \frac{p_A(1-p_AN_x)}{1-p_AN_xN_y} - \frac{(1-p_A)[1-(1-p_A)N_x]}{1-(1-p_A)N_xN_y}\right].$$

IC (A1) and (A3) can then be rewritten as: $\xi(1,y) \leq 0$ and $\psi(x,1) \leq 0$. Note that $\xi(x,y)$ strictly decreases in α and $\psi(x,y)$ strictly decreases in β. Thus for any given y, a cutoff $\alpha(y)$ exists such that $\xi(0,y) > 0$ if and only if $\alpha < \alpha(y)$. If $\alpha < \alpha(y)$, then biased A's agenda pushing benefit strictly exceeds his reputation cost, and thus his best response is $x^{BR}(y) = 0$. A similar cutoff $\beta(x)$ can be defined for any x such that if $\beta < \beta(x)$, biased B's best response is $y^{BR}(x) = 0$.

Next, note that $\xi(x,y)$ strictly increases in x and $\psi(x,y)$ strictly increases in y. Hence if $\alpha \geq \alpha(y)$, $\xi(0,y) < 0$. Thus there exists a unique $x' \in (0,1)$ such that $\xi(x',y) = 0$. Also, for any given x, if $\beta > \beta(x)$, $\psi(x,0) < 0$. Thus there exists a unique $y' \in (0,1)$ such that $\psi(x,y') = 0$. If α and β are sufficiently large, $\xi(x,y) = 0$ implicitly defines the best response function of biased A to y such that: $x^{BR}(y) \in (0,1)$ for all $y \in [0,1]$. Similarly, $\psi(x,y) = 0$ implicitly defines biased B's best response function such that: $y^{BR}(x) \in (0,1)$ for all $x \in [0,1]$. Because $\xi(x,y)$ and $\psi(x,y)$ are continuous in x and y, both $x^{BR}(y)$ and $y^{BR}(x)$ are continuous.

Let ξ_1, ξ_2 be the partial derivative of ξ with respect to x and y; ψ_1, ψ_2 are similarly defined. Then, $d/dy\ x^{BR}(y) = -\xi_2/\xi_1$, and $d/dx\ y^{BR}(x) = -\psi_1/\psi_2$. From the analysis above, $\xi_1 > 0$, $\psi_2 > 0$. Moreover, differentiate with respect to N_y, we have:

$$\xi_2 = \frac{(2p_A-1)N_x}{(2-N_xN_y)^2} - \alpha\theta_A(1 - N_x)\left[\frac{p_A^2}{(1-p_AN_xN_y)^2} + \frac{(1-p_A)^2}{(1-(1-p_A)N_xN_y))^2}\right].$$

Substitute in ξ(x,y) $= 0$, ξ_2 is equal to $a\theta_A(1 - N_x)$ times:

$$\frac{p_A}{1 - p_A N_x N_y}\left[\frac{1}{2 - N_x N_y} - \frac{p_A}{1 - p_A N_x N_y}\right]$$

$$+ \frac{1 - p_A}{1 - (1 - p_A)N_x N_y}\left[\frac{1}{2 - N_x N_y} - \frac{1 - p_A}{1 - (1 - p_A)N_x N_y}\right],$$

which is negative. Given $s_A = 0$, biased A knows that he receives $\pi_{A(1,0)}$ with a higher probability p_A, thus the first part of the above expression, which is negative, dominates. Similarly, $\psi_1 < 0$. Thus the best responses of biased A and B, if positive, are strictly increasing.

PROOF OF PROPOSITION 1:

There are three possible types of agenda pushing equilibrium. First, an interior equilibrium in which biased A and B report truthfully with positive probability: $x > 0$, $y > 0$, and $\psi(x,y) = 0$ and $\xi(x,y) = 0$. Second, a corner equilibrium in which biased A and B always report 1: $x = y = 0$, and $\psi(0,0) > 0$ and $\xi(0,0) > 0$. Third, a hybrid equilibrium in which one biased agent always reports 1, and the other reports truthfully with positive probability: $x = 0, y > 0$; or $x > 0, y = 0$.

To show that an interior agenda pushing equilibrium exists if α, β are sufficiently high, recall from Lemma 2 that in this case $\psi(0,0) < 0$ and $\xi(0,0) < 0$. Thus $y^{BR}(0) > 0$, $y^{BR}(1) < 1$, and similarly, $x^{BR}(0) > 0$, $x^{BR}(1) < 1$. Also, $x^{BR}(y)$ and $y^{BR}(x)$ are continuous and strictly increasing if positive. Note that $\psi\ (x,y) = 0$ implicitly defines a function $y = h(x)$, which is continuous and strictly increasing, and $h(0) < 0$, $h(1) > 1$. Because $y^{BR}(0) - h(0) > 0$, $y^{BR}(1) - h(1) < 0$, by the intermediate value theorem, there exists some $x \in (0,1)$ such that $y^{BR}(x) - h(x) = 0$. This implies that $\psi(x, y^{BR}(x)) = 0$, thus the two best response functions must intersect at x, $y > 0$.

Moreover, in equilibrium, it can be shown that $\xi_1\psi_2 - \xi_2\psi_1 > 0$: whenever biased A and B's best responses intersect, biased A's best response function has a steeper slope than biased B's. Also, since the best responses are strictly increasing, the interior agenda pushing equilibrium is unique. A sufficient condition for an interior equilibrium to exist is $\alpha > \alpha^i$, $\beta > \beta^i$, where:

$$\alpha^i \equiv \frac{(2p_A - 1)(1 - p_A\theta_A\theta_B)(1 - (1 - p_A)\theta_A\theta_B)}{(1 - \theta_A)(2 - \theta_A\theta_B)(1 - 2p_A(1 - p_A)\theta_A\theta_B)};$$

$$\beta^i \equiv \frac{(2p_A - 1)(1 - p_A\theta_A\theta_B)(1 - (1 - p_A)\theta_A\theta_B)}{(1 - \theta_B)(2 - \theta_A\theta_B)(1 - 2p_A(1 - p_A)\theta_A\theta_B)}.$$

These cutoff values are defined such that $\xi(0,0) = 0$, $\psi(0,0) = 0$ respectively at α^i, β^i.

A unique corner equilibrium in which biased agents always report $m_A = 1$ and $m_B = 1$ exists if $\alpha \leq \alpha^i$ and $\beta \leq \beta^i$. In this case, $\xi(0,0) \geq 0$ and $\psi(0,0) \geq 0$, hence $x^{BR}(0) = 0$, $y^{BR}(0) = 0$. By continuity, $x^{BR}(y) = 0$ for y sufficiently close to 0 and $y^{BR}(x) = 0$ for x sufficiently close to 0. Also, since $x^{BR}(1) < 1$, $y^{BR}(1) < 1$, if there were equilibria in which $x > 0$, $y > 0$, the best response functions cannot intersect only once, violating the property that biased A's best response is always steeper than biased B's in equilibrium. Thus $x = 0$, $y = 0$ is the only equilibrium.

Finally, if $\alpha \leq \alpha^i$, but $\beta > \beta^i$, then $\xi(0,0) \geq 0$, $\psi(0,0) < 0$. From the above analysis, a unique $y' \in (0,1)$ exists such that $\psi(0,y') = 0$. There are two possibilities: (i) if $\xi(0,y') \geq 0$, then in the unique equilibrium, biased A always reports $m_A = 1$ while biased B reports $m_B = 0$ with probability y' if $m_A = 0$. (ii) If $\xi(0, y') < 0$, then in the unique equilibrium, biased A and B both report truthfully with positive probabilities. Which equilibrium may occur depends on parameter values. Similarly, if $\xi(0,0) < 0$, $\psi(0,0) \geq 0$, then in equilibrium, either $x > 0$, $y = 0$ (if $\xi(x',0) = 0$, $\psi(x',0) \geq 0$) or $x > 0$, $y > 0$ (if $\xi(x',0) = 0$, $\psi(x',0) < 0$).

PROOF OF COROLLARY 1:
See Web Appendix A.

PROOF OF PROPOSITION 2:
Biased A's incentive constraints are given in the proof of Lemma 1 with $y = w = 1$. For any x, z, the difference in the RHS of biased A's IC (A3) and IC (A4),

$$\alpha(2p_A - 1)[\pi_{A(0,0)} + \pi_{A(1,1)} - \pi_{A(0,1)} - \pi_{A(1,0)}],$$

is positive. Thus if IC (A3) holds strictly or with equality, IC (A4) must hold strictly. The only possibilities are $x \in [0,1)$, $z = 1$ and $x = 1$, $z \in [0,1)$. Moreover, the LHS of IC (A3) is positive. If $x = 1$, $z \in [0,1)$, the RHS of IC (A3) is negative while the LHS is positive, a contradiction. As shown in Lemma A1, both biased A and B need to randomize for an uninformative equilibrium to exist, which is impossible here. Hence the equilibrium must be an agenda pushing one.

Next, recall from the text that biased A reports $m_A = 0$ if $s_A = 0$ with probability x^d. Note that IC (A3) never holds if biased A's weight on reputation $\alpha < \alpha^d$, where the cutoff

$$\alpha^d \equiv \frac{(2p_A - 1)(1 - p_A\theta_A)[1 - (1 - p_A)\theta_A]}{(2 - \theta_A)(1 - \theta_A)(1 - 2p_A(1 - p_A)\theta_A)}$$

increases in θ_A. In this case, biased A always reports $m_A = 1$. If $\alpha \geq \alpha^d$, the LHS of IC (A3) is smaller than the RHS at $x^d = 0$. Because the LHS strictly increases in x^d and the LHS strictly decreases in it, there exists a unique x^d such that IC (A3) holds with equality, where x^d is defined by:

$$\text{(A7)} \qquad \frac{2p_A - 1}{2 - N_x^d} - \alpha\theta_A\left[\frac{1}{N_x^d} - \frac{p_A(1 - p_A)}{1 - p_A N_x^d} - \frac{p_A(1 - p_A)}{1 - (1 - p_A)N_x^d}\right] = 0.$$

PROOF OF COROLLARY 2:
See Web Appendix A.

PROOF OF PROPOSITION 3:
See Web Appendix A.

PROOF OF COROLLARY 4:
See Web Appendix A.

PROOF OF PROPOSITION 4:
See Web Appendix A.

PROOF OF PROPOSITION 5:
See Web Appendix A.

PROOF OF PROPOSITION 6:
See Web Appendix A.

STRATEGIC OBJECTIVE AGENT AND CHANNEL CHOICE:
See Web Appendix B.

REFERENCES

Arrow, Kenneth J. 1985. "Informational Structure of the Firm." *American Economic Review*, 75(2): 303–07.

Arrow, Kenneth J., and Roy Radner. 1979. "Allocation of Resources in Large Teams." *Econometrica*, 47(2): 361–85.

Austen-Smith, David. 1990. "Information Transmission in Debate." *American Journal of Political Science*, 34(1): 124–52.

Benabou, Roland, and Guy Laroque. 1992. "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility." *Quarterly Journal of Economics*, 107(3): 921–58.

Blume, Andreas, Oliver J. Board, and Kohei Kawamura. 2007. "Noisy Talk." *Theoretical Economics*, 2(4): 395–440.

Chevalier, Judith, and Glenn Ellison. 1999. "Career Concerns of Mutual Fund Managers." *Quarterly Journal of Economics*, 114(2): 389–432.

Crawford, Vincent P. 2003. "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions." *American Economic Review*, 93(1): 133–49.

Crawford, Vincent P., and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica*, 50(6): 1431–51.

Dekel, Eddie, and Michel Piccione. 2000. "Sequential Voting Procedures in Symmetric Binary Elections." *Journal of Political Economy*, 108(1): 34–55.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel. 2003. "Persuasion Bias, Social Influence, and Unidimensional Opinions." *Quarterly Journal of Economics*, 118(3): 909–68.

Dewatripont, Mathias, and Jean Tirole. 1999. "Advocates." *Journal of Political Economy*, 107(1): 1–39.

Evans, John H., R. Lynn Hannan, Ranjani Krishnan, and Donald V. Moser. 2001. "Honesty in Managerial Reporting." *Accounting Review*, 76(4): 537–59.

Goltsman, Maria, Johannes Horner, Gregory Pavlov, and Francesco Squintani. 2009. "Mediation, Arbitration and Negotiation." *Journal of Economic Theory*, 144(4): 1397–1420.

Hendricks, Ken, and R. Preston McAfee. 2006. "Feints." *Journal of Economics and Management Strategy*, 15(2): 431–56.

Ivanov, Maxim. Forthcoming. "Communication via a Strategic Mediator." *Journal of Economic Theory*.

Kawamura, Kohei. 2006. "Anonymity, Equal Treatment, and Overconfidence: Constraints on Communication May Enhance Information Transmission." Unpublished.

Krishna, Vijay, and John Morgan. 2004. "The Art of Conversation: Eliciting Information from Experts through Multi-stage Communication." *Journal of Economic Theory*, 117(2): 147–79.

Li, Hao, Sherwin Rosen, and Wing Suen. 2001. "Conflicts and Common Interests in Committees." *American Economic Review*, 91(5): 1478–97.

Li, Ming, and Tymofiy Mylovanov. 2008. "Credibility for Sale: The Effect of Disclosure on Information Acquisition and Transmission." Unpublished.

Li, Wei. 2007. "Changing One's Mind When the Facts Change: Incentives of Experts and the Design of Reporting Protocols." *Review of Economic Studies*, 74(4): 1175–94.

Li, Wei. Forthcoming. "Well Informed Intermediaries in Strategic Communication." *Economic Inquiry*.

Lim, Terence. 2001. "Rationality and Analysts' Forecast Bias." *Journal of Finance*, 56(1): 369–85.

Malueg, David A., and Shunichi O. Tsutsui. 1996. "Duopoly Information Exchange: The Case of Unknown Slope." *International Journal of Industrial Organization*, 14(1): 119–36.

Mirman, Leonard J., Larry Samuelson, and Edward E. Schlee. 1994. "Strategic Information Manipulation in Duopolies." *Journal of Economic Theory*, 62(2): 363–84.

Morris, Stephen. 2001. "Political Correctness." *Journal of Political Economy*, 109(2): 231–65.

Ottaviani, Marco, and Peter N. Sorensen. 2006. "Professional Advice." *Journal of Economic Theory*, 126(1): 120–42.

Prat, Andrea. 2005. "The Wrong Kind of Transparency." *American Economic Review*, 95(3): 862–77.

Prendergast, Canice, and Lars Stole. 1996. "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning." *Journal of Political Economy*, 104(6): 1105–34.

Scharfstein, David S., and Jeremy C. Stein. 1990. "Herd Behavior and Investment." *American Economic Review*, 80(3): 465–79.

Shapiro, Carl. 1986. "Exchange of Cost Information in Oligopoly." *Review of Economic Studies*, 53(3): 433–46.

Sobel, Joel. 1985. "A Theory of Credibility." *Review of Economic Studies*, 52(4): 557–73.