Peddling Influence through Intermediaries: Propaganda*

Wei Li
University of California, Riverside
wei.li@ucr.edu.

July, 2007

Abstract

Information may be transmitted directly from a sender to a receiver, or indirectly through intermediaries. How do intermediaries affect the reporting truthfulness of an informed sender? When does he prefer using intermediaries? An objective sender or intermediary always passes on information truthfully, while a biased one wants to push a particular agenda but also has reputational concerns. This paper shows that intermediaries reduce an agenda-pushing sender's reputation cost, but they also lessen his influence on the receiver. Biased agents' truth-telling incentives are strategic complements, and each additional intermediary reduces everyone's reporting truthfulness. If the sender's prior reputation is sufficiently good and his signal sufficiently informative, ex ante, he prefers using intermediaries. If the sender's prior reputation is sufficiently low and his reputational payoffs sufficiently convex, he prefers direct communication.

JEL classification: C70, D72, D82, D83

Keywords: strategic communication through intermediaries, indirect communication, agenda pushing, reputational concerns.

*I am grateful for the insightful comments from Abhijit Banerjee, Glenn Ellison, Mathias Dewatripont, Botond Kőszegi, Hao Li, R. Preston McAfee, Edward Schlee, Jean Tirole, and the seminar participants at MIT, Kellogg School of Management, UCR, ASU, USC, UCSB, Academia Sinica, Taiwan, the 9th Southwest Economic Theory conference, the 7th Canadian Economic Theory Conference and the 2007 Summer Meeting of the North American Econometric Society.

1 Introduction

Suppose a campaign manager of a political candidate is interested in discrediting an opponent as well as appearing objective in the eyes of voters. He may launch a direct advertisement attacking the opponent. Under the Bipartisan Campaign Reform Act (BCRA) enacted in 2002, he must disclose his identity.[1] Alternatively, he may convey the information to another political organization or an activist group who may choose what to tell the voters. Such groups, for example, the 527 organizations or Internet forums, are frequently not subject to the same disclosure rules.[2] Should the campaign manager run a negative advertisement directly or use another political organization?

In a similar vein, a government administration intent on pushing a particular agenda or selling a policy may present the relevant information to the public directly. However, doing so may be risky, especially if the agenda is unsupported by 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).[3] The media then chooses what to inform the public. The public's reactions have major policy ramifications. Such practices are common, for instance, information such as prewar intelligence on Iraq was intentionally leaked to news media (CNN 2006a, CNN 2006b); the recent trial and conviction of I. Lewis Libby Jr. indicated that classified intelligence was disclosed to reporters for political purposes (Lewis 2007). What are the advantages and drawbacks of influencing public opinion though intermediaries?

Many papers have studied the incentives of a biased sender who aims to influence the action of a receiver by manipulating the information he sends (Crawford and Sobel 1982, Dewatripont and Tirole 1999, Chevalier and Ellison 1999, Morris 2001, Prat 2005, Ottaviani and Sorensen 2006, among others).[4] The informed sender in these papers may be biased out of reputational concerns, or because he has a specific agenda, but he always tries to influence the receiver directly. Clearly, as the opening examples

[1] Political candidates for federal office 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."

[2] The 527 groups are tax-exempt organizations that engage in political advocacy. They are not regulated by the Federal Election Commission and may raise unlimited amount of soft money contributions. In the 2006 election cycle, for example, the Democratic/liberal 527 groups spent over $45 million and the Republican/conservative ones spent over $64 million. The data was based on IRS records released on February 28, 2007. For more details, see http://www.opensecrets.org/527s/.

[3] In news media, anonymity is widely granted, but this practice is currently under debate. For instance, in the first week of April 2005, 47% of all A-section articles published in the New York Times used anonymous sources, 46% of which were identified as "officials" or "aids" only (Okrent 2005).

[4] Throughout this paper, the sender(s) of information is male and the decisionmaker is female.

suggest, information often travels indirectly, through intermediaries.

This paper develops a model of communication through strategic intermediaries. An intermediary receives a message from an informed sender, then sends a message of his own to an uninformed decision-maker. Such an intermediary has two effects on a sender. First, his presence affects the sender's reputation. He and the sender can each be objective or biased: an objective agent is assumed to pass on what he hears, but a biased one wants to sell a particular agenda *and* to appear objective. The decisionmaker, then, rationally blames both the intermediary and the sender if she receives an agenda-pushing message that contradicts the evidence, which she learns after making a decision. This blame sharing effect reduces the reputation cost of distorting information. Second, the presence of an intermediary affects the credibility of the sender's message. In my model, to avoid any information aggregation complications, the intermediary is assumed to have no (or very little) information of his own. This implies that the decisionmaker is less likely to believe in what she hears from an intermediary–who may only introduce further distortion—than directly from the sender. This makes the sender's agenda pushing less effective.

The relative magnitude of these two effects depends on how the sender and the intermediary's truth-telling incentives interact. The main insight emerging from this model is that their truth-telling incentives are strategic complements. For a biased sender, there are two countervailing effects if the intermediary is slightly more truthful. On the one hand, because they share the blame for an agenda-pushing message, the sender pays a higher reputation cost, making it less attractive for him to lie. On the other hand, the final message becomes more credible, making it more attractive for him to lie. Surprisingly, the net effect is unambiguous: the sender also reports more truthfully.

To see why, observe that in this model, the decisionmaker acts first and learns the truth later. This difference in available information is crucial. When the decisionmaker hears from the intermediary, she believes that it is accurate with some probability because of the presence of objective agents. Thus the message still has a major effect on her despite the possible distortions. Afterwards, she observes the true state, at which point a wrong message is more likely to result from distortion than from a wrong signal of nature. Because she attributes, ex post, a larger share of any agenda-pushing message to the agents' distortions, the intermediary shares the sender's blame more than reduces his influence. More truthful reporting from the intermediary thus increases the sender's reputation cost proportionally more than it increases his agenda pushing effectiveness.

Because of this complementarity, a biased sender always lies more with intermediaries than without. For the decisionmaker, a direct message—viewed as a message transmitted through a completely truthful intermediary—is more credible than an indirect one. The voting public in the opening examples, then, should evaluate anything learned from intermediaries cautiously: not only the political activist groups or the media may introduce bias of their own, they also worsen the sender's truth-telling incentives. This very complementarity, though, may also aid the decisionmaker in reducing information loss from indirect communication. Each biased agent's truth-telling incentives are shown to increase in how much *any* agent cares about his reputation. Thus if the decisionmaker cannot control all agents, perhaps for legal or practical reason, she can still improve everyone's truth telling by making it more costly for the intermediary to lie. This suggests that policies such as stricter enforcement of disclosure laws or higher standards for anonymity-granting make everyone more truthful.

This paper provides a further insight on the ex ante choice of a biased sender. It shows that whether a sender prefers a particular communication channel hinges on how important the saving in his reputation cost from using intermediaries is relative to the loss of effectiveness in his agenda pushing. A biased sender prefers direct communication for two reasons. First, his reputational concerns may be so low that they are strictly dominated by the loss in influence. Second, and somewhat more subtlely, a sender with a biased image prefers direct communication if he needs to appear highly objective, for instance, to get re-elected. In this case, any agenda pushing message, delivered with or without intermediaries, does little to boost his biased image. Instead, he is better off sending a direct message not associated with his agenda: it is a good signal of his objectivity, and it is more effective. That is, direct communication may be used more often if an administration needs to drastically improve the public's perception of its objectivity.

In contrast, the sender prefers indirect communication if his information is highly informative and he has moderately high reputational concerns. If he chooses direct communication, then in the event that his information does not support his agenda, he risks losing (almost) all reputation if he lies: he knows that the message is likely wrong. And he can ill afford it due to his reputational concerns. Moreover, his high signal quality implies that he still exerts a lot of influence even with intermediaries. Thus a government expecting highly accurate information may nonetheless choose to hide behind intermediaries.

Interestingly, the above analysis suggests a new rationale for media bias. The intermediary here cannot influence the public at all without the sender. In certain situations, a biased sender can be shown to prefer direct communication to an intermediary of sterling objectivity. Intuitively, a more biased intermediary is more attractive to the sender because he shoulders more blame after a wrong message. Thus, the intermediary may cultivate a biased image to gain access to information it would not have otherwise.

Several recent papers extend the Crawford and Sobel (1982) framework to more general communication protocols, allowing some role for non-strategic intermediaries. Blume, Board, and Kawamura (2007) add garbling to the communication process such that, instead of the sender's message, a random message may reach the receiver with some probability. They show that more information may be transmitted with noise than that is possible in Crawford and Sobel, partly because the noise dampens the receiver's response to any message and thus reduces the sender's incentive to distort his signal. Goltsman, Horner, Pavlov, and Squintani (2007) show that if there is a neutral, unbiased intermediary (mediator) who adds the noise optimally, communication protocols employed by Blume, Board, and Kawamura (2007) may yield the highest ex ante payoff for the receiver. The current paper, instead, focuses on strategic intermediaries with reputational concerns. Here, intermediaries not only introduce distortions (noise), they may also worsen the truth-telling incentives of all biased agents and thus reduce the message quality significantly.

This paper is most related, in term of the setup of the model without intermediaries, to Bénabou and Laroque (1992) in which the objective type is assumed to report honestly, but the biased type (insiders) have reputational concerns: they need to appear credible in order to manipulate the market's belief of an asset effectively. Morris (2001) considers the case where the objective agent also faces reputational concerns. He shows that there exists a "politically correct" equilibrium in which the message associated with bias may be avoided by an objective agent sufficiently concerned about future reputation. Both assume direct communication. In this model, the objective agent is behavioral: how strategic, biased agents behave across communication channels is the main question of interest.

Section 2 sets up the indirect communication game. Section 3 and 4 analyze, respectively, the biased sender's behavior without and with an intermediary. Section 5 studies a sender's ex ante choice of communication channels. Section 6 extends the model, discusses several main assumptions, and Section 7 concludes. All proofs are collected in the Appendix.

2 The Indirect Communication Game: Setup

A decision needs to be made based on the true state of the world $\eta \in \{0, 1\}$. Each state occurs with equal probability. There are three agents: A, B and C. Agent C is the decisionmaker who needs to take an action $a \in \Re$. Her optimal action is to choose a equal to the probability she attaches to $\eta = 1$. In the government administration example, for instance, the true state may be "no military threat" (state 0) and "high military threat" (state 1). Decisionmaker C, then, represents the voters who need to choose an appropriate level of war mobilization.

Agent A is the government. He, and only he, observes a private signal s_A about the state of the world. This signal is equal to the true state with probability $p_A > \frac{1}{2}$; otherwise it is wrong. Agent B, the media, is assumed to be a pure intermediary who has no signal of his own. This assumption simplifies away the information aggregation complications, and makes it possible to focus on how A's incentives to report truthfully depends on the communication channel. It also captures the situation that A's signal, perhaps of a classified nature, is significantly more informative than that of B's.[5] After observing his signal, A sends a message $m_A \in \{0, 1\}$ to B, who in turn sends a message $m_B \in \{0, 1\}$ to C. Information flows only in one direction, from A to B to C. Each agent can only observe the message sent to him directly. Moreover, the true state and all messages are assumed to be observable but unverifiable, thus no transfers can be made based on the messages.

Agent i ($i = A, B$) may be either objective (type o) or biased (type b). Each agent's type is independently drawn from $\{o, b\}$: $Pr(i = o) = \theta_i$, $Pr(i = b) = 1 - \theta_i$. Parameter θ_i is referred to as the prior objectivity of agent i in this paper. An objective agent is assumed to report his information (s_A or m_A) honestly. Honesty here is interpreted either as an institutional goal or a behavioral trait. Some media and non-profit organizations may adhere to an ethical standard of only informing the public in an impartial way; some people may simply prefer behaving honestly, as suggested by psychological experiments (Evans, Hannan, Krishnan, and Moser 2001).[6]

[5] Main results of this paper hold qualitatively if B observes a sufficiently uninformative signal s_B, e.g., $Pr(s_B = \eta) = p_B \approx \frac{1}{2}$. In a companion piece, Li (2006) considers the case where the intermediary is a well informed expert in the market for credence goods. See further discussions on well-informed intermediaries in Section 6.

[6] For instance, BBC's editorial guideline states that "We will be objective and even handed 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."

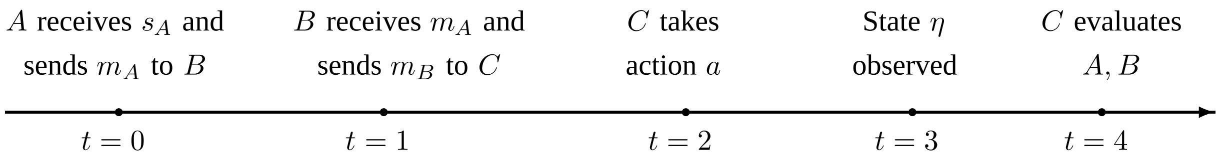

Figure 1: Timeline of the Indirect Communication Game

A biased agent always favors action $a = 1$, but he also wants to appear objective due to reputational concerns. Denote agent i's posterior probability of being objective as π_i, which is formed after C observes the true state η. Biased A and B's payoffs are assumed to be, respectively:

$$u_A = a + \alpha V_A(\pi_A) \quad \text{and} \quad u_B = a + \beta V_B(\pi_B).$$

The first half of biased i's payoff function is C's action. The more likely C takes action $a = 1$, which is the favorite agenda of a biased agent, the better off he is. The second half is a reduced form formulation capturing a biased agent's reputational payoffs: V_i depends on agent i's posterior objectivity, and $\alpha, \beta \in [0, \infty)$ are the weights A and B attach to their reputations. In summary, the indirect communication game is illustrated in Figure 1.

In this paper, V_i is assumed to be biased i's posterior objectivity: $V_i = \pi_i$. This reduced form formulation, used in many existing papers, reflects the fact that an agent is less effective in influencing the decisionmaker if he is considered highly biased (Scharfstein and Stein 1990, Prendergast and Stole 1996, Ottaviani and Sorensen 2006). In a dynamic setting, the agent's reputational payoffs are determined by C's decision problem in the future, and the linear form is not without loss of generality, as pointed out by Prat (2005), Ottaviani and Sorensen (2006) and shown in Li (2007). The following three examples illustrate why in many settings the biased sender and intermediary may be concerned about their perceived objectivity, which may be either linear or convex. For this reason, the ensuing analysis also highlights certain implications if the agent's reputational payoffs are convex.

Example 1: Midterm elections. Continue with the government example, where C is the voting public and A is the government who may have real evidence supporting a war, or have no evidence, but a pro-war agenda. Suppose that the public has acted and then observed the true state. Afterwards, the administration faces a midterm congressional election. Here C needs to determine what control A's party should be given over war related policy. She takes action $a_2 \in \Re$ to minimize $(a_2 - \pi_A)^2$, and her optimal action is to set $a_2 = \pi_A$. That is, A and his party's control over the war (measured by its number of seats in the Congress) depends linearly on the public's perception of whether A is biased. $\square$

Example 2: Future policy. Similar to Morris (2001), suppose that a biased sender cares about his reputation only because he wants to influence the decisionmaker in the second stage. Suppose the second stage game is identical to that in the first, but with a new (and independent) state η_2, a new noisy signal with precision p_A, and a new action $a_2 \in \Re$ chosen by the decisionmaker to minimize loss function $(a_2 - \eta_2)^2$. Because this is the last stage, biased A always reports $m_A = 1$. The decisionmaker chooses $a_2 = Pr(\eta_2 = 1|m_B)$, and biased A's reputational payoff in the first stage becomes $V_A(\pi_A) = \frac{2p_A - 1}{2 - \pi_A} + 1 - p_A$, which is increasing and convex in π_A. $\square$

Example 3: Media bias. Suppose that B is a cable news channel with a possible pro-war bias; C remains the public. The first stage is as described above. In the second stage, the public needs to decide whether to stop the war right away ($a_2 = 0$) or to continue ($a_2 = 1$). If C continues the war, its outcomes depend on the true state of the world η_3, which is ex ante good or bad ($\eta_3 = \{g, b\}$) with equal probability. It is simplest to equate the outcome with the state: it is either good ($g > 0$) or bad ($b < 0$). C receives zero if she stops the war. At the beginning of the second stage, B receives an informative signal s_B: $Pr(s_B = \eta_2) = p_B$. As this is the end of the game, biased B always reports that the war will go well ($\eta_3 = g$). If C hears a pro-war report from B, her expected payoff of continuing the war is simply $0.5(g + b) + (p_B - 0.5)(g - b)\pi_B$.

Next, if $g + b \geq 0$, B's value of information increases linearly in π_B. In this case, C always chooses $a_2 = 1$, but the more objective B is, the more C is willing to pay for his news (in the form of subscriptions). If $g + b < 0$, however, C's default action is to stop. This implies that B's value of information is 0 if π_B is below a cutoff value, and increasing and affine in π_B otherwise.[7] Intuitively, if B appears very biased at the beginning of the second stage, his news has no value: the decisionmaker always stops the war and never subscribes to B. If B appears sufficiently objective, the value of his news services increases in his posterior objectivity. Hence B's first stage reputational payoff is increasing and piece-wise linear. $\square$

In this game, biased agent i simply sends a message $m_i \in \{0, 1\}$ as a function of his information (s_A or m_A). Given message m_B, C chooses an action a. Later, she rationally updates her opinion on A and

[7] The value of B's report is 0 if $\pi_B < -\frac{0.5(g+b)}{(p_B - 0.5)(g-b)}$. Otherwise, it is $(p_B - 0.5)(g - b)\pi_B + 0.5(g + b)$.

B's objectivity π_A, π_B as a function of their prior objectivity, the message received and the observed state. This paper looks for perfect Bayesian equilibrium (PBE): each agent chooses a message to maximize his expected payoff, given his information, the other agent's strategy as well as C's action and inferences.

Although messages are assumed to be unverifiable, they are not cheap talk in this model. Due to the presence of objective type who always passes on information truthfully, any message is informative and useful to C. This also implies that if a biased agent lies to push his agenda, his message is more likely to be wrong than that from an objective agent, and he pays a higher reputation cost than reporting truthfully. Consequently, there are no babbling equilibria in which the message is uncorrelated with the agent's signal, and ignored by the decisionmaker.

3 The Baseline Case: Direct Communication

This section examines the case where A sends a message to C directly.[8] It illustrates the basic tradeoff a biased source faces in a simple setting, and thus serves as a useful benchmark against which the indirect communication model will be compared. Also, it is relevant when voters or consumers need to evaluate platforms or advertisements directly from potentially biased sources.

Objective A simply reports $m_A = s_A$, but biased A wants to push his agenda $\eta = 1$ and to appear objective. Given signal s_A, biased A chooses m_A to maximize his expected payoff:

$$EU_A(m_A|s_A) = Pr(\eta = 1|m_A) + \alpha E_\eta[Pr(A = o|m_A, \eta)|s_A].$$

The first part is the decisionmaker's optimal action upon receiving A's message, and the second part is A's (expected) posterior objectivity multiplied by the weight he attaches to his reputation, where the expectation is taken with respect to state η. Clearly, A's message choice is determined by the net difference in his expected payoff of sending $m_A = 1$ versus sending $m_A = 0$: $EU_A(m_A = 1|s_A) - EU_A(m_A = 0|s_A)$.

Before analyzing a biased sender's behavior, it is helpful to begin by identifying the key equilibrium properties in this model. The following definition greatly eases the exposition:

[8] Formally, this is equivalent to the case where A hires a known objective intermediary B ($\theta_B = 1$), or if the biased B faces an infinitely high reputation cost: $\beta = \infty$.

Definition 1 *An "agenda-pushing equilibrium" is an equilibrium in which a biased agent reports* 1 *truthfully if his information supports his agenda* ($s_A = 1$ *or* $m_A = 1$)*; and reports* 0 *truthfully with a probability strictly smaller than* 1 *if it does not support his agenda* ($s_A = 0$ *or* $m_A = 0$)*.*

The corresponding strategy is referred to as an "agenda-pushing strategy". Then it can be shown that:

Lemma 1 *Every equilibrium of this game is an agenda-pushing equilibrium.*

This result shows that honesty is never the best policy for a biased agent. If it were an equilibrium for A to report $s_A = 0$ truthfully, decisionmaker C knows that everyone reports truthfully ($m_A = s_A$) and acts accordingly. Then by reporting $m_A = 1$, A can exert the maximum amount of influence on C at no reputation cost. Hence biased A strictly gains from deviating and lying if $s_A = 0$, which is a contradiction.

Lemma 1 also shows that it never pays for a biased agent to intentionally distance himself away from message $m_A = 1$. More precisely, even though the direction of bias is known, there does not exist a perverse equilibrium in which A reports $s_A = 0$ truthfully, but lies after receiving $s_A = 1$. If there were such an equilibrium, $m_A = 1$ indicates $s_A = 1$ for sure and is thus very convincing. Moreover, $m_A = 1$ becomes a better sign of objectivity. This encourages A to deviate and report $m_A = 1$ if $s_A = 0$, which is impossible under the putative equilibrium.

Given Lemma 1, biased A's strategy can be restricted to an agenda-pushing one. Let him report $m_A = 0$ with probability x^d if $s_A = 0$. Then the difference in his expected utility if he reports $m_A = 1$ versus $m_A = 0$ can be decomposed into two parts: how strongly A's message changes C's action, and how much it changes his posterior objectivity. First, examine A's agenda pushing effectiveness. The net difference in C's action induced by A's message is, for both signals:

$$Pr(\eta = 1|m_A = 1) - Pr(\eta = 1|m_A = 0) = \frac{p_A + (1 - p_A)(1 - x^d)(1 - \theta_A)}{1 + (1 - x^d)(1 - \theta_A)} - (1 - p_A).$$

The second part is true because given A's strategy, the true signal is $s_A = 0$ if $m_A = 0$. This difference is strictly positive due to the presence of objective A. It also increases in x^d because the more truthful a biased A is, the more credible $m_A = 1$ becomes, and the more C believes it.

Second, examine A's reputation cost. Given signal s_A, if A reports $m_A = 1$ instead of $m_A = 0$, the net difference in his posterior objectivity is:

$$\alpha[Pr(A = o|m_A = 0) - E_\eta[Pr(A = o|m_A = 1, \eta)|s_A]].$$

This difference is non-negative because $m_A = 0$ is more likely to come from an objective agent. In fact, A always suffers a loss in reputation by sending $m_A = 1$. Moreover, this difference is decreasing in x^d, because the more truthful A is, the less C modifies her view of his objectivity from the message itself. Intuitively, if even a biased agent reports very truthfully, $m_A = 0$ is not a strong signal of objectivity; nor is $m_A = 1$ a strong signal of bias.

A biased sender lies against signal $s_A = 0$ if his net gain from agenda pushing by sending $m_A = 1$ outweighs the net loss in his expected reputational payoff. The following proposition summarizes A's behavior in direct communication:

Proposition 1 (Direct Communication) *There exists a cutoff value $\overline{\alpha}$ such that, in the unique agenda-pushing equilibrium, A always reports $m_A = 1$ if $\alpha \leq \overline{\alpha}$; and reports $s_A = 0$ truthfully with probability $x^d > 0$ if $\alpha \geq \overline{\alpha}$.*

Proposition 1 shows that if $s_A = 0$, a biased agent reports truthfully with a positive probability if he attaches a sufficiently high weight to his future reputation.[9] A natural question, then, is how A's reporting accuracy x^d depends on his own characteristics, such as his prior objectivity and signal quality. In the context of the opening examples, one may ask whether the political 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 following result provides some answers.

Corollary 1 *(1) Reporting accuracy and A's prior objectivity. If A's reputation weight $\alpha \leq \frac{1}{2}$, A always lies. If $\alpha > \frac{1}{2}$, then given signal quality p_A, there exist cutoff values $\theta_A^1, \theta_A^2 \in (0,1)$ such that x^d increases in θ_A if $\theta_A \in [0, \theta_A^1]$; decreases if $\theta_A \in [\theta_A^1, \theta_A^2]$; and becomes zero if $\theta_A \in [\theta_A^2, 1]$.*

(2) Reporting accuracy and A's signal quality. Given α and θ_A, if $\alpha\theta_A \geq \frac{1}{2}$, then x^d first decreases in his signal quality p_A; but eventually increases as p_A becomes sufficiently high. If $\alpha\theta_A < \frac{1}{2}$, then x^d first decreases in p_A and becomes 0 as p_A becomes sufficiently high.

Corollary 1 shows that reporting accuracy tends to be non-monotonic in a biased source's characteristics. First, A is most truthful when he has the most to gain or lose in term of reputation, which occurs when his prior objectivity is in the intermediate range. For example, when a political candidate is relatively

[9] The cutoff values $\overline{\alpha}$ and the equilibrium truthful reporting probability x^d are defined in the proof in the Appendix.

unknown. To see this, observe that if biased A is, ex ante, considered to be very objective ($\theta_A > \theta_A^2$), a wrong message has a minimal impact on his reputation, because C believes that A is most likely objective but the signal was wrong. At the other extreme ($\theta_A \approx 0$), even though A lies almost completely, he reports more truthfully as θ_A increases. Because $m_A = 0$ is almost a sure sign of objectivity, a marginal increase in truthful reporting makes him appears very objective. As θ_A increases further, A's reputation gain from reporting $m_A = 0$ decreases; and his reputation cost from reporting $m_A = 1$ falls. Therefore biased A's reporting accuracy actually falls as θ_A becomes sufficiently high.[10]

Second, a biased source may lie more, not less, as his signal becomes more accurate. A's signal quality has two opposing effects on his truth telling. Suppose that A faces sufficiently high reputational concerns such that $\alpha\theta_A > \frac{1}{2}$.[11] On the one hand, the more accurate his signal is, the more informative it becomes. Thus a positive message has a stronger impact on C's action, increasing his incentive to lie. On the other hand, as p_A increases, whenever message $m_A = 1$ turns out wrong, it is more likely that A has lied. This higher reputation cost decreases A's incentive to lie. Corollary 1 shows that if the signal is very uninformative ($p_A \approx \frac{1}{2}$), even an objective agent is often wrong, thus A's gain in agenda pushing dominates, and he lies more. However, when p_A becomes sufficiently high, a wrong message is (almost) a sure sign of bias. Lying leads to a complete loss of reputation, which outweighs any gain from agenda pushing, and he lies less eventually.

Next, Example 2 and 3 show that A may have a convex reputational payoff function because information about his objectivity itself has positive value in the future. Politicians and media outlets perceived to be very objective may exert disproportionately more influence on the decisionmaker than those with uncertain objectivity. To study how this may affect a biased source's truth telling, consider a marginal increase in the convexity of A's reputational payoff. Suppose that $V_A(\pi_A) = (\pi_A)^\rho,\ \rho > 1$, then:

Corollary 2 *If in equilibrium $x^d > 0$, and ρ is sufficiently close to 1, then x^d increases in ρ if θ_A is sufficiently close to 0; but decreases in ρ if θ_A is sufficiently close to θ_A^2.*

Recall that A is always better off reporting $m_A = 0$ in term of his reputation: $m_A = 1$ is costly

[10] 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 very high or very low. Because they are interested in long term reputation formation, the agent's truth-telling incentives when he has an intermediate prior reputation may be ambiguous.

[11] If this condition is satisfied, either A attaches a sufficiently high weight to his reputation, or he has a sufficiently high objectivity that his reputation loss is significant if he lies.

because of its association with bias. Corollary 2 shows that two new effects surface if a better reputation becomes disproportionately more attractive. First, the "top prize" from $m_A = 0$, which does not vary with the (later) observed state, becomes relatively more valuable. Second, the expected reputational payoff from lying increases as well because it is a risky gamble. Therefore, biased A's net reputation cost may increase or decrease in ρ, depending on his prior objectivity. If A is very likely biased, he can boost his posterior objectivity significantly by reporting $m_A = 0$ more often, while that from reporting $m_A = 1$ remains approximately zero. The first effect dominates and he becomes more truthful. In contrast, an agent with a very objective image has little to gain from reporting $m_A = 0$. Thus the payoff from reporting more truthfully increases less than that in his expected reputational payoff from reporting $m_A = 1$. As a result, the second effect dominates and the agent with a high prior objectivity actually reports less truthfully.

This simple corollary suggests that if higher levels of perceived objectivity matter more in the future, a source perceived to be very biased becomes more honest, e.g. a think tank with clear ideological bias may push its agenda less often. However, if the source has a good prior, e.g. a major news outlet, this may encourage and reward further distortion: if he is lucky and the distorted message turns out right, he becomes very credible in the future; and his reputation cost is relatively low if he is wrong. As a result, A may become less "fair and balanced".

4 Indirect Communication

Building on the direct communication model, this section shows how using a possibly biased intermediary affects the effectiveness of A's message and his perceived objectivity. It also considers the information loss associated with indirect communication, and examines its implications in the media and in law.

4.1 Equilibrium with an intermediary

Objective A, B pass on information truthfully. Similar to Lemma 1, it can be shown that every equilibrium is an agenda-pushing equilibrium. Thus biased A, B adopt an agenda-pushing strategy such that they report $s_A = 0$ and $m_A = 0$ with probability x, y respectively. Given this strategy, biased B chooses a message m_B to maximize his expected payoff: $EU_B(m_B|m_A) = Pr(\eta = 1|m_B) + \beta E_\eta[Pr(B = o|m_B, \eta)|m_A]$.

The difference in C's induced action if B reports $m_B = 1$ instead of $m_B = 0$ is:

$$Pr(\eta = 1|m_B = 1) - Pr(\eta = 1|m_B = 0) = \frac{p_A - 0.5}{\underbrace{0.5[1 + (1-\theta_A)(1-x)]}_{Pr(m_A=1)} + \underbrace{0.5(\theta_A + (1-\theta_A)x)(1-\theta_B)(1-y)}_{Pr(m_A=0,\ m_B=1)}}.$$

This difference is always positive: the presence of objective agents implies that B's message always pushes C's belief in the direction of the message. In fact, even if all biased agents lie completely, C still believes more in $\eta = 1$ if $m_B = 1$. Formally, $Pr(\eta = 1|m_B = 1) > Pr(\eta = 0|m_B = 1)$ if $x = y = 0$.[12] Moreover, this difference increases in the truth-telling probabilities x, y: the more truthful A and B are, the more likely C is swayed by B's message. Note that holding A's behavior constant ($x = x^d$), this net influence on C is smaller than that in direct communication. That is, a potentially biased message from B is less effective in agenda pushing than that from A. Since A's signal is the only available information, all that B can do is to induce further distortion, as can been seen from the part labeled $Pr(m_A = 0, m_B = 1)$ in the expression above.

Biased A's expected payoff is similar to that in the direct communication, except that now he needs to anticipate B's message. At first sight, it may seem unclear how this uncertainty affects A: a truthful message of $m_A = 0$ may still be distorted by B, which affects C's action and A's perceived objectivity. Interestingly, both A's net benefit from agenda pushing and his net reputation cost of lying are multiplied by a common factor: $Pr(m_B = 0|m_A = 0)$, the probability that B passes on A's message 0. Specifically, the net benefit for A to report $m_A = 1$ versus $m_A = 0$ is:

$$Pr(m_B = 0|m_A = 0)[Pr(\eta = 1|m_B = 1) - Pr(\eta = 1|m_B = 0)].$$

Similarly, A's net reputation cost is $\alpha Pr(m_B = 0|m_A = 0)[Pr(A = o|m_B = 0) - E_\eta[Pr(A = o|m_B = 1, \eta)]|s_A]]$. The reason is that A knows that $m_A = 1$ will be passed on for sure; but if $m_A = 0$, then $m_B = 1$ reaches C with probability $Pr(m_B = 1|m_A = 0)$. The pivotal event for agent A—which drives his message choice—is whether he could change what C hears. His message only matters when it does. Intuitively, because C cannot observe m_A, both A's influence on C and his posterior reputation are filtered through B's message. This observation greatly simplifies the analysis, because A's incentive to lie vis-a-vis B's can be analyzed with this factor taken out. Thus, A and B receive the same benefit from

[12] This also shows that, by reporting $m_B = m_A$, the objective B passes on the most accurate information he has.

agenda pushing *relative* to his reputation cost. Any difference in their reporting accuracy must be driven by differences in A and B's reputation costs.

How does B's message affect both agents' reputation costs? In comparison with the direct communication case, C's evaluation of A and B's objectivity becomes more subtle because she does not observe m_A. If $m_B = 0$, C knows that A's signal is $s_A = 0$ for sure; neither A nor B has distorted it. However, three things may have occurred if she hears $m_B = 1$: the true signal $s_A = 1$; agent B passes on a distorted message $m_A = 1$, or B has distorted A's message to push his agenda. Biased A, B weigh the benefit from agenda pushing against their own reputation cost. The following proposition characterizes the equilibrium of the indirect communication game:

Proposition 2 *(2.1) If both agents place sufficiently low weights on their reputations (α and β sufficiently close to 0), or if their prior objectivities θ_A and θ_B are sufficiently high, there exists a unique agenda-pushing equilibrium in which they lie completely: $x = 0, y = 0$.*

(2.2) If both agents place sufficiently high weights on their reputations ($\alpha \geq \tilde{\alpha}$ and $\beta > \tilde{\beta}$), there exists a unique agenda-pushing equilibrium in which $x, y \in (0, 1)$.

First, Proposition 2 shows that if an agent's prior objectivity is sufficiently high, the agent's message has little marginal impact on his reputation. In this case, he gains strictly from agenda pushing and lies completely. The same is true if a biased agent has little concern for appearing objective. It also shows that biased agents with moderate reputational concerns—high enough so that one cannot afford to lie completely even if the other agent does so—report information that does not support their agenda truthfully with positive probability.[13]

Second, the key to understand biased A, B's truth-telling incentives is to see how the reporting accuracy of one affects that of the other. Suppose that B is slightly more truthful (y increases slightly), two opposing effects surface. On the one hand, B's message becomes more credible, thus the decisionmaker is more likely to choose $a = 1$. As a result, A's agenda pushing effectiveness increases in B's truth telling, encouraging A to lie more (x falls). On the other hand, A now faces a higher reputation cost if $m_B = 1$, because C rationally attributes more blame of initiating a biased message to A. This encourages A to

[13] The cutoff values $\tilde{\alpha}$ and $\tilde{\beta}$ are defined in the appendix. Intuitively, this condition guarantees that each agent is unwilling to lie completely, regardless of the other agent's behavior. Also, the cutoff value for A in direct communication $\bar{\alpha}$ is smaller than that $\tilde{\alpha}$. This means that A needs to be more concerned about his reputation to report truthfully with an intermediary.

lie less (x rises). Intuitively, A and B free ride on each other: each agent's net reputation cost increases in the other's truthful reporting, but decreases in his own. The relative magnitude of these two opposing effects, then, determines whether A becomes more truthful or not.

As a simple illustration, suppose $p_A = 1$, which means that if the message is wrong ($m_B \neq \eta$), C knows that someone has lied. After a wrong message, then, A is considered objective with probability $Pr(A = o|m_B = 1, \eta = 0)$, which increases in x but decreases in y. If B reports more truthfully, A's net reputation cost increases by:

$$\frac{\Gamma}{0.5[1 - (\theta_A + (1 - \theta_A)x)(\theta_B + (1 - \theta_B)y)]}, \tag{1}$$

where Γ is a positive expression defined in the appendix. Observe that the denominator is $Pr(m_B = 1, \eta = 0)$, how likely the wrong message $m_B = 1$ reaches the decisionmaker. Next, A's agenda-pushing effectiveness, $Pr(\eta = 1|m_B = 1) - Pr(\eta = 1|m_B = 0)$, increases in y by the following amount:

$$\frac{\Gamma}{0.5[2 - (\theta_A + (1 - \theta_A)x)(\theta_B + (1 - \theta_B)y)]}. \tag{2}$$

Here, the denominator is the *total* probability that a message $m_B = 1$ is received. Because $Pr(m_B = 1)$ is clearly larger than $Pr(m_B = 1, \eta = 0)$, Expression (1) is bigger than Expression (2). Therefore at the same level of reporting truthfulness, if B is more truthful, A's reputation cost increases more than his gain from the additional credibility B's message now has. Consequently, it becomes more costly for him to lie and he wants to report more truthfully. In equilibrium, A and B's truth telling are strategic complements: x and y increase together.

Intuitively, this is because C chooses her action and forms her belief about the agent's objectivity based on different information. For example, in a political campaign, negative information about a candidate may reach the voters and affect their decision before its truth is revealed.[14] When C hears m_B, she assigns a high probability to the true signal being positive than to A and B's lying. Thus the message still has a major effect on her despite the possible distortions. Ex post, however, C knows for sure that either A or B has lied. Therefore, B shares the sender's blame more than reduces his influence.

[14] For instance, in the Bush-McCain campaign of 2000, South Carolina voters were asked, in an anonymous push-poll, "Would you be more likely or less likely to vote for John McCain for president if you knew he had fathered an illegitimate black child?". Later, John McCain lost South Carolina, effectively ending his run for the presidency. It turned out that McCain has an adopted Bangladeshi daughter with whom he campaigned.(CNN 2000)

Third, the biased agent's relative reporting accuracy depends on how much they care about their reputation. First, if A and B are symmetric, they lie with the same probability: $x = y$ if $\theta_A = \theta_B$ and $\alpha = \beta$. Moreover, no asymmetric equilibrium in which biased A and B behave differently exists. Suppose in equilibrium, $x > y$, then conditional on a wrong signal from nature, C is more likely to attribute message $m_B = 1$ to B than to A. Also, B pays a higher reputation cost of *not* reporting $m_A = 0$ than A. This leads to an impossibility: A and B receive the same (relative) benefit in term of agenda pushing, but B pays a higher reputation cost than A by fabricating $m_B = 1$. Second, given similar prior objectivity, a biased agent more concerned about his reputation reports more truthfully in equilibrium: $x > y$ if $\alpha > \beta$; $x < y$ if $\alpha < \beta$. Third, it can be shown that if one biased agent, say agent A places a sufficiently low weight on his reputation ($\alpha \approx 0$), but the other agent, agent B, is very concerned about his reputation ($\beta > \tilde{\beta}$), then in the unique equilibrium, A lies completely and B reports truthfully sometimes.

This complementarity between biased agents suggests that changes in one agent's reputational concerns affect the other agent. This is of practical importance for a decisionmaker who cannot influence all agents directly, perhaps due to existing anonymity granting rule of the media or the laws protecting whistle blowers.[15] She may consider imposing a higher cost-financial or reputational-on the intermediary. For instance, several courts in the recent years have grappled with setting the appropriate legal guidelines for when intermediaries such as Internet service providers can be compelled to divulge the identities of their customers.[16] The following shows how a policy change such as stronger protection of customer identity in the law or stricter rules of sourcing in the media, modeled as an increase in β, may influence all.

Proposition 3 (Overall effect of disciplining the intermediary) *Suppose that $x, y > 0$ in the agenda-pushing equilibrium of the indirect communication game. Then biased A and B become more (less) truthful if either agent becomes more (less) concerned with their reputation: x, y increase in α, β. Moreover, if $\theta_A \approx \theta_B$, a biased agent responds more to any increase in his own reputational concerns than that in the other: $\frac{\partial x}{\partial \beta} < \frac{\partial y}{\partial \beta}$, $\frac{\partial x}{\partial \alpha} > \frac{\partial y}{\partial \alpha}$.*

[15] For an example, see U.S. Environmental Protection Agency's rules on "Confidentiality, Anonymity, & Whistleblower Protection" at http://www.epa.gov/oig/ombudsman-hotline/protection.htm.

[16] In Doe v. Cahill, 884 A.2d 451 (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 about a politician on the Internet. In Melvin v. Doe, 836 A.2d 42 (Pa. 2003), the Pennsylvania Supreme Court has indicated that a test balancing the First Amendment right to anonymous speech with the right to address unprotected speech should be developed.

Clearly, as a media outlet, B reports more truthfully if he faces higher fines for granting anonymity too casually. Proposition 3 shows that this makes it more costly for A to lie as well. Therefore the decisionmaker can improve the overall reporting accuracy by increasing the reputation cost of the intermediary. 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".[17]

However, the reverse is also true: information deteriorates quickly even if only one agent cares less about his reputation, e.g., a politician whose public life is drawing to an end.[18] Moreover, the positive effect of such policies may be quite limited. The reason is that a biased source responds more to a change in his own reputational concerns than that in an intermediary's (if they have similar prior objectivities). For example, the media may become more scrupulous in reporting due to stricter anonymity granting rules, but the source barely reduces his lying. Since it takes only one biased agent to distort the information, a wrong message may still reach C with a high probability.

5 Comparing Communication Channels

This section addresses two questions. The first question is how the decisionmaker interprets messages across communication channels. For instance, how should the voters evaluate a piece of news if it comes from a political opponent directly or if it comes from an activist group citing confidential or obscure sources? The second question is how a biased source ranks the communication channels for propaganda purposes. Which channel makes him better off?

5.1 Truthful Reporting in Different Channels

Sometimes an agent may only communicate in a particular way. In elections, buyers of political advertisements are required to disclose their identities. In government and corporations, officials may not be allowed to release information to the public directly. This subsection compares a biased source's reporting

[17] In a June 23, 2005 memo titled "Assuring Our Credibility" by Bill Keller, executive editor of the New York Times.

[18] Moreover, if agent A lies completely in equilibrium, a small change in B's reputation cost does not affect A's behavior. A sufficiently large increase in β is necessary for A to report more truthfully.

accuracy in these two channels.

Proposition 1 and 2 show that a biased agent reports $s_A = 0$ truthfully with probability x^d without intermediary B, and x with him. The intermediary renders A less effective in agenda pushing, but also lowers A's reputation cost. The net effect, though, is unambiguous:

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

This result, a consequence of Proposition 3, shows that B saves more in A's reputation cost than reduces A's influence. This is because direct communication is equivalent to indirect communication where the biased intermediary has infinitely high reputational concerns ($\beta = \infty$). Proposition 3 shows that if the intermediary becomes less concerned with his reputation, which is the case when A changes from direct to indirect communication, both agents report less truthfully.[19]

As an illustration, Figure 2 shows that if $p_A = 0.9$ and A and B are symmetric ($\alpha = \beta = 1, \theta_A = \theta_B = \theta$), the probability A reports truthfully with intermediary, x, always lies below x^d, his truth-telling probability in direct communication.

Figure 2: Reporting Accuracy With and Without Intermediary
P_A=0.9

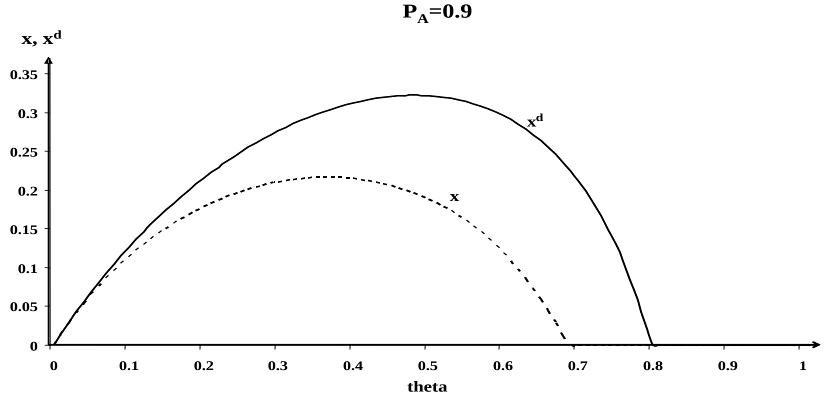

[19] The insight that the intermediary reduces A's reputation cost more than his agenda pushing effectiveness also holds when there are many intermediaries, which is presented as an extension in Section 6.

More importantly, this shows that not only an intermediary may introduce bias, he also reduces all other biased agents' truth-telling incentives. The decisionmaker, then, should be cautious in evaluating messages learned indirectly. Information learned directly tends to be significantly more credible than that learned though more agents. The decisionmaker, then, prefers direct communication because the associated information loss is smaller.[20]

Observe that indirect communication reduces A's reputation cost primarily when he lies and reports $m_A = 1$. Therefore, A's behavior may change if such saving is negligible. Suppose that A has to appear highly objective to exert influence in the future, then a direct message from A may be *less* accurate than the message he sends to B.

Corollary 4 *Biased A lies more in direct communication, $x^d \leq x$, if θ_A is sufficiently low and his reputational concerns are sufficiently convex in his posterior objectivity.*

The key intuition here is that indirect communication is of little use when a biased source needs to drastically improve the public's perception of him. Because of A's low prior objectivity, C thinks of him as very biased if she receives a positive message, whether or not intermediary B is involved. Also, since only high levels of posterior objectivity matter for biased A, the small saving from lying through B is negligible. In comparison with direct communication, A exerts less influence but faces a similar net reputation cost, thus he lies less with intermediary B.[21]

5.2 Choosing communication channels

Which communication channel should a biased source choose *before* the signal is observed? In this section, biased A commits to a channel that leads to a higher expected payoff, not knowing whether his agenda is supported by evidence or not.[22] A government can commit to indirect communication by classifying

[20] In fact, indirect communication leads to two types of information losses. First, it is prone to the propagation of distorted information and second, it makes true signal $s_A = 1$ much less useful for the decisionmaker. Using measures such as that of mean absolute error, it can be shown that the information loss is increasing and concave in the fraction of biased agents if there are very few of them.

[21] This result suggests that, intriguingly, if A reports very accurately through B, the decisionmaker C may receive a more accurate report from indirect communication, depending on the characteristics of B.

[22] This assumption is widely used in the literature on information sharing among oligopolists, where information exchange decisions are taken prior to the arrival of private information (such as the realization of cost). Thus issues of incentives to reveal information truthfully when a firm is already aware of its own costs are not considered. See, for instance, Shapiro (1986), Malueg

certain information for security reasons.

Let $EU_A^d(s_A), EU_A^i(s_A)$ denote, respectively, A's expected equilibrium payoff with direct and indirect communication given signal s_A. Because the state is distributed symmetrically, the difference in his expected payoff is simply $\frac{1}{2}\sum_{s_A}[EU_A^d(s_A) - EU_A^i(s_A)]$. Biased A's channel choice is not obvious: indirect communication saves reputation cost, but a direct message is more credible, especially because A reports more truthfully in direct communication. The following proposition gives some conditions under which A prefers direct communication; and conditions under which A prefers using an intermediary.

Proposition 4 (Channel Choice with Commitment) *(4.1) Suppose that in equilibrium, $x^d > 0, x = 0$. Then A prefers indirect communication.*

(4.2) Suppose that in equilibrium, $x^d = x = 0$. Then A prefers direct communication if his weight on reputation α is sufficiently close to 0; or if his prior objectivity θ_A is sufficiently close to 1; or if B's weight on reputation β is sufficiently high. But A prefers indirect communication if α is sufficiently close to α, his signal is sufficiently informative, and β is sufficiently low.

(4.3) Suppose that in equilibrium, $x^d > 0, x > 0$. Then A is ex ante indifferent.

First, Proposition 4 shows that if A reports truthfully sometimes in the direct channel, but lies completely with intermediaries, he prefers indirect communication. This occurs if A has moderately high reputational concerns, or if his prior objectivity is moderately high. The reason is that, by using intermediary B, A's saving in reputation cost is so high that he can "afford" to lie completely while he cannot do so directly.

Second, if A attaches little weight to his reputation, he prefers direct communication because his reputation loss in either channel is negligible. Similarly, if A has an excellent prior objectivity, the correctness of his message has a minimal impact on his reputation. In both cases, the marginally better reputation he gains from using intermediary B is outweighed by the loss in agenda pushing. More interestingly, an intermediary perceived to be rather biased may be more helpful to a biased source than a more objective one. Proposition 4 shows that if A can afford to lie completely in both channels, he chooses direct communication if B has very high reputational concerns and is thus very truthful. This is

and Tsutsui (1996) and the references within. If there is no commitment, then biased A chooses a communication channel after knowing his signal, and the receivers, both agent B and the decisionmaker C make additional inferences about the signal from A's channel choice. See Section 6 for further discussions on the case without commitment.

because C still attributes almost all blame to A if the final message is wrong. Instead, A chooses indirect communication if his signal is very accurate, and intermediary B has very little reputational concerns. Such an intermediary is very apt to lie, therefore shoulders a big share of the blame if the final message is wrong. In this case, the saving in A's reputation cost is sufficiently high for him to use the intermediary. For this reason, a government with little reputation to lose may nonetheless go through a likely biased intermediary if his information is highly accurate.

Relatedly, these results suggest an interesting role for biased intermediaries. Because agent B has no private signal, he cannot influence C at all without source A. Biased B, then, may prefer to appear biased: doing so encourages a biased source to communicate through him, and in turn makes him more influential. That is, even if the public prefers objective media outlets, some may still cultivate a less objective image in order to gain access to sensitive information.

Third, if biased A reports truthfully with some probability in both channels, he is ex ante indifferent. The reason is that a biased agent's payoff function amounts to a weighted sum of C's posterior beliefs (of the true state and of A's type). Thus A's ex ante expected payoff of sending any message is equal to the sum of C's priors. The channel A has chosen, however, does lead to different payoffs once the signal is observed. In particular, if $s_A = 0$, his payoff is higher if he has committed to indirect communication. Recall from Corollary 3 that A lies more with intermediary B ($x^d > x$) because intermediary B saves more in A's reputation cost than reduces A's effectiveness. Thus A receives a higher expected payoff using B. If $s_A = 1$, his payoff is higher if he has committed to direct communication. This is because $m_A = 1$ is likely to be correct, and A is unlikely to receive the lowest reputation, namely $m_A = 1, \eta = 0$.

More generally, A's preferred communication channel depends on his reputational payoff, which may be convex in his posterior objectivity. A's channel choice, then, depends on how a convex reputational payoff interacts with his own characteristics. The following result suggests, however, biased A is no longer indifferent between the communication channels even if he reports truthfully with positive probabilities with a linear reputational payoff. Again, suppose that $V(\pi_A) = (\pi_A)^\rho,\ \rho > 1$.

Corollary 5 *(1) If ρ is sufficiently high, A prefers direct communication if he has sufficiently low prior objectivity θ_A. (2) If ρ is sufficiently close to 1, A prefers indirect communication if his signal is sufficiently informative, and θ_A is close to θ_A^2.*

First, a highly convex reputational payoff makes direct communication more attractive for a likely biased A. Corollary 4 shows that indirect communication does not help here because low levels of objectivity matter little in A's reputational payoffs. By sending a direct message, not only A becomes more persuasive now, it also easier for him to obtain the best reputation later: $m_A = 0$ is a much stronger signal of objectivity.

Second, a convex reputational payoff makes indirect communication more attractive for a likely objective A whose signal is very informative. Similar to Proposition 4, A still faces a severe loss of reputation if his message turns out to be wrong in direct communication. Moreover, a small increase in ρ makes the risky action of sending a positive message more attractive than before. Thus A's net reputation cost falls, and he is better off using intermediary B.

6 Discussions and Extensions

This section discusses several main assumptions on the number of intermediaries, the sender's interim choice of communication channels, as well as intermediary's lack of private information. It also suggests how the agents' behavior may be affected if these assumptions were varied.

A. Multiple intermediaries. This model can be easily extended to the case with many intermediaries. Suppose that there are k agents, each of whom can send a message to his immediate successor. Agent $k+1$ is the decisionmaker. For simplicity, let $p_A = 1, \theta_i = \theta, \alpha_i = \alpha, i = 1, 2, ..., k$. Namely, the agents are symmetric and the first agent, agent $i = 1$, receives a perfect signal: $Pr(s_1 = \eta) = 1$. All other assumptions remain. Then the following can be shown:

Proposition 5 *If α is sufficiently low, or θ is sufficiently close to 1, all agents report $m_i = 1$. If $\alpha(1-\theta) > \frac{1-\theta^k}{2-\theta^k}$, then the unique equilibrium is an agenda-pushing one in which biased i reports $m_i = 0$ with the same probability $x^k \in (0, 1)$. Moreover, each agent lies more as the number of agents increases: x^k decreases in k.*

This result shows that if the biased agents are sufficiently concerned about their reputation, they report truthfully with probability x^k. But each agent's reporting accuracy decreases in the number of intermediaries. To see why no asymmetric equilibrium exists, note that if the last message is correct

($m_k = \eta$), all agents are considered equally objective. The posterior probability each agent is objective is θ if $\eta = 1$ and $\frac{\theta}{\theta+(1-\theta)x^k}$ if $\eta = 0$. However, if the last message is wrong, e.g. $m_k = 1, \eta = 0$, then some biased agent i has lied (all agents after him follow his message $m_{i+1} = 1$). For the decisionmaker, if the distortion occurs before or after i, i's reputation is clearly unaffected. The pivotal case is if $m_{i-1} = 0$, but $m_i = 1$. For biased i, all that matters is how likely he can change the final report from $m_k = 0$ by reporting truthfully (and if all other agents pass it on truthfully) to $m_k = 1$. Similar to the two agents case, it can be shown that each agent receives the same net benefit from agenda pushing relative to his net reputation cost. Thus one's position does not matter when indirect communication is used: each biased agent's truth-telling incentives depend on his prior objectivity and the weight he attaches to his reputation, not where he is.[23] Also, because all agent's truth-telling incentives are strategic complements, each additional intermediary reduces every biased agent's reputation cost of lying more than he reduces their agenda-pushing effectiveness.

B. Interim communication channel choice. This paper focuses on biased A's channel choice before the signal is observed. Proposition 4 shows that, even though A is ex ante indifferent between the channels, his interim payoffs are different. A natural question arises: if there is no commitment, what channel should A choose given his signal?[24] For instance, should the government always inform the public directly if the intelligence supports its pro-war agenda?

To analyze this, assume that an objective agent is still behavioral, thus he is equally likely to use either direct or indirect channel. But a biased agent, given his signal s_A, chooses both a channel and a message m_A in that channel. All other assumptions remain. Then it can be shown that first, similar to the first part of Proposition 4, if A has extremely low reputational concerns such that $\theta_A \approx 1$ and $\alpha \approx 0$, he prefers direct communication regardless of the signal. Because in this case, not only a biased source cares little about his lower reputation cost from indirect communication, his agenda pushing in direct communication is also very effective due to his high priors.

[23] In a different context, a similar pivotal argument have been used in Li, Rosen, and Suen (2001) and Dekel and Piccione (2000) to show that the order of voting does not affect the voting outcomes in equilibrium. In a non-strategic social network context, DeMarzo, Vayanos, and Zwiebel (2003) shows that one's influence on other people in a social network depends not only on his information accuracy, but also on his position in a given social network. In their model, the agents report truthfully, but have a "persuasion bias". Namely they fail to account for possible repetitions in the information that reaches them.

[24] A's channel choice is also one way for him to manipulate the informativeness of his signal, in particular the decisionmaker C's belief about his objectivity. Strategic manipulation of signal informativeness in a duopoly context is examined by Mirman, Samuelson, and Schlee (1994).

Second, suppose that a biased source is moderately concerned about his reputation and there exists a mixed strategy equilibrium in direct communication. Then there does not exist a pooling equilbrium such that biased A always prefers one channel to the other. That is, a government with moderate reputational concerns will not use one channel exclusively, direct or indirect, for all possible news that he received. To see this, imagine that a biased source eschews direct communication, then only objective A uses it. Then, a message delivered directly must reflect the true signal and come from an objective agent. Thus biased A is strictly better off sending $m_A = 1$ directly. Similarly, A will not avoid indirect communication. Also, it seems unlikely for a biased source to use direct communication if $s_A = 1$ and indirect communication if $s_A = 0$. If so, a positive message from indirect communication is very ineffective and A is likely to deviate to direct communication. Because a message of $m_A = 1$ is very credible since it shows that the signal is positive for sure; and a message of $m_A = 0$ shows one is objective. This suggests that a biased source may use both channels with positive probability even if the signal does not support his agenda.

C. Better informed intermediaries. The current model examines the case of a pure intermediary. In many marketing, medicine and lobbying settings, however, B may in fact have good information of his own. Li (2006) considers the case where an objective expert 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 decisionmaker. Because of B's superior information, he would, if objective, always follow his own signal, which means that independence becomes the prized sign of the objective. It also means that biased A and B are no longer sharing blame. B cannot use the excuse that he is misled: he is not supposed to be influenced in the first place. One main implication is that A and B's lying are now substitutes: any improvement in B's incentive worsens that of A's. In fact, this effect may be so strong that the net effect is negative. A practical consequence is that improving ethical standard or campaign finance laws may lead, perversely, to more biased information being transmitted.

7 Conclusion

Should the government in the opening example send information through intermediaries? An intermediary would share the blame of disseminating false information, thereby reducing the government's reputation cost. But indirect messages are less effective. The government should inform the public directly if it has extremely low reputation concerns; or if it needs to appear highly objective despite an existing biased image. It should leak to intermediaries if its signal is very informative, and he has a good image in the eyes of the voters.

The intermediary may benefit from appearing biased, because it encourages a biased source to send him information he would otherwise not receive. For the voters, though, the intermediary reduces eveyone's incentive to report truthfully, in addition to introducing potential bias of its own. The voters may improve the overall reporting accuracy by increasing the financial or legal costs of the intermediary, but this positive effect may be very limited.

Indirect communication matters in studying the structure of firms and organizations. Previous literature analyzed the information aggregation role of the firm and the resulting optimal firm structure (Arrow 1985, Arrow and Radner 1979). In many firms, however, both formal, direct communication and informal, indirect communication exist. How the co-existence of various communication channels affects a firm's information processing as well as the associated incentive problems is a topic of further research. This model may also be extended to study military operations. For instance, false intelligence has been fed to unfriendly media by Pentagon as part of a disinformation campaign (Shanker and Schmitt 2004). Using intermediaries may be particularly fruitful here: typical military operations are zero sum, thus direct communication is unlikely to be effective (Crawford 2003, Hendricks and McAfee 2006).

Appendix

Proof of Lemma 1:

Recall that the decisionmaker C's belief is simply $a = Pr(\eta = 1|m_A)$. To simplify notations, denote C's action after receiving m_A as $a_1^A \equiv Pr(\eta = 1|m_A = 1)$ and $a_0^A \equiv Pr(\eta = 1|m_A = 0)$ respectively. Then, for biased A to report his signals truthfully, the following two incentive constraints need to hold:

$$EU_A(m_A = 1|s_A = 1) \geq EU_A(m_A = 0|s_A = 1) \tag{3}$$

$$\Rightarrow \quad a_1^A - a_0^A \geq \alpha E_\eta[Pr(A = o|m_A = 0, \eta) - Pr(A = o|m_A = 1, \eta)|s_A = 1];$$

$$EU_A(m_A = 1|s_A = 0) \leq EU_A(m_A = 0|s_A = 0) \tag{4}$$

$$\Rightarrow \quad a_1^A - a_0^A \leq \alpha E_\eta[Pr(A = o|m_A = 0, \eta) - Pr(A = o|m_A = 1, \eta)|s_A = 0].$$

We now analyze these two incentive constraints. Suppose that a biased A reports $s_A = 0$ truthfully with probability x, and reports $s_A = 1$ with probability z.

Claim 1: there does not exist a fully mixed equilibrium in which $x, z \in (0, 1)$. If this claim were false, biased A lies with positive probability after both signals. Then the LHS of IC (3) and IC (4) are clearly equal; the difference in their RHS can be shown to be:

$$\begin{aligned} & \alpha(2p_A - 1)[Pr(A = o|m_A = 0, \eta = 1) - Pr(A = o|m_A = 0, \eta = 0)] \\ + \quad & \alpha(2p_A - 1)[Pr(A = o|m_A = 1, \eta = 0) - Pr(A = o|m_A = 1, \eta = 1)] \leq 0. \end{aligned}$$

This inequality is true because A's signal is assumed to be informative: $p_A > \frac{1}{2}$, and that, given this strategy, a wrong report in either direction is always a worse sign of one's objectivity, who always reports truthfully. Thus IC (3) and IC (4) cannot both bind, which is a contradiction.

Claim 2: there does not exist a truth-telling equilibrium in which $x = z = 1$. If this claim is false, then the LHS of IC (3) and IC (4) become $2p_A - 1 > 0$. The RHS of both incentive constraints become zero, because if all agents report truthfully, the posterior objectivity is simply the prior, which does not depend on the message or the observed state. This clearly violates IC (4). This shows that a biased agent will never be completely truthful.

Claim 3: there does not exist an equilibrium in which $x = 1, z \in [0, 1)$. Suppose such an equilibrium exists, then it must be that IC (3) binds or fails to hold, and IC (4) holds strictly. Given this strategy, the LHS of A's incentive constraint is:

$$\begin{aligned} & a_1^A - a_0^A \\ = \quad & Pr(\eta = 1|s_A = 1) - \frac{Pr(m_A = 0|\eta = 1)Pr(\eta = 1)}{Pr(m_A = 0|\eta = 1)Pr(\eta = 1) + Pr(m_A = 0|\eta = 0)Pr(\eta = 0)} \\ = \quad & \frac{2p_A - 1}{1 + (1 - \theta_A)(1 - z)} > 0. \end{aligned} \tag{5}$$

Simple algebra can show that the RHS of the incentive constraints are strictly negative, therefore it is impossible for IC (4) to hold, a contradiction. This shows that no perverse equilibrium in which the biased agent distances himself away from message $m_A = 1$ if $s_A = 1$.

Finally, the only remaining possibility is $x \in [0,1), z = 1$, which is precisely the agenda-pushing equilibrium defined in the text. □

Proof of Proposition 1:

Recall from the text that if $s_A = 0$, A reports $m_A = s_A$ truthfully with probability x^d. Then A's posterior objectivity is given by Bayes' rule as follows:

$$Pr(A = o|m_A = 0, \eta = 0) = Pr(A = o|m_A = 0, \eta = 1) = \frac{\theta_A}{\theta_A + (1-\theta_A)x^d};$$
$$Pr(A = o|m_A = 1, \eta = 0) = \frac{(1-p_A)\theta_A}{1 - p_A + p_A(1-\theta_A)(1-x^d)};$$
$$Pr(A = o|m_A = 1, \eta = 1) = \frac{p_A\theta_A}{p_A + (1-p_A)(1-\theta_A)(1-x^d)}.$$

A's net gain of reporting $m_A = 1$ over reporting $m_A = 0$ is:

$$\begin{aligned} & a_1^A - a_0^A \\ = \; & \frac{Pr(m_A = 1|\eta = 1)Pr(\eta = 1)}{Pr(m_A = 1|\eta = 1)Pr(\eta = 1) + Pr(m_A = 1|\eta = 0)Pr(\eta = 0)} - Pr(\eta = 1|s_A = 0) \\ = \; & \frac{2p_A - 1}{1 + (1-\theta_A)(1-x^d)} > 0. \end{aligned} \tag{6}$$

This gain is strictly increasing in x^d: the more honest biased A is, the more informative a message of $m_A = 1$ becomes, and the more likely C believes $\eta = 1$. On the other hand, if $s_A = 0$, A's net reputation cost of reporting $m_A = 1$ over $m_A = 0$ is:

$$\alpha[Pr(A = o|m_A = 0, \eta) - (1-p_A)Pr(A = o|m_A = 1, \eta = 1) - p_A Pr(A = o|m_A = 1, \eta = 0)]. \tag{7}$$

This cost is strictly decreasing in x^d: the higher x^d is, the more honest A is, and $m_A = 1$ is less likely a sign of bias. If at $x^d = 0$, the LHS of IC (4) is strictly larger than the RHS, then IC (4) never holds and IC (3) holds strictly. This occurs when A's weight on reputation $\alpha < \overline{\alpha}$, where

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

In this case, in the unique equilibrium, biased A always reports $m_A = 1$ regardless of his signal.

If $\alpha \geq \overline{\alpha}$, the LHS of IC (4) is smaller than the RHS at $x^d = 0$. Because the LHS is strictly increasing in x^d and the LHS is strictly decreasing, there exists a unique truth-telling probability such that IC (4) binds. The equilibrium mixing probability is implicitly defined by:

$$\frac{2p_A - 1}{2-(\theta_A+(1-\theta_A)x^d)} = \frac{\alpha\theta_A(1-\theta_A)(1-x^d)[1-2p_A(1-p_A)(\theta_A+(1-\theta_A)x^d)]}{[\theta_A+(1-\theta_A)x^d][1-p_A(\theta_A+(1-\theta_A)x^d)][1-(1-p_A)(\theta_A+(1-\theta_A)x^d)]}. \tag{8}$$

Moreover, because IC (3) holds strictly if IC (4) binds, in the unique equilibrium, biased A reports $s_A = 1$ truthfully, but reports $s_A = 0$ truthfully with probability x^d. □

Proof of Corollary 1:

Proposition 1 shows that biased A reports $s_A = 0$ truthfully with $x^d \in [0, 1)$ in the unique equilibrium. Note that the cutoff value $\overline{\alpha}$ decreases in θ_A. At $\theta_A = 0$, $\overline{\alpha} = p_A - \frac{1}{2}$. If $\alpha \leq p_A - \frac{1}{2}$, A always lies: $x^d = 0$ regardless of θ_A. For any $\alpha \geq \frac{1}{2}$, there exists a cutoff value θ_A^2 such that $g(p_A, \theta_A^2, \alpha) = 0$, where

$$g(p_A, \theta_A, \alpha) \equiv \frac{2p_A - 1}{2-\theta_A} - \frac{\alpha(1-\theta_A)[1-2p_A(1-p_A)\theta_A]}{[1-p_A\theta_A][1-(1-p_A)\theta_A]}.$$

That is, if $\theta_A > \theta_A^2$, $x^d = 0$; and $x^d > 0$ otherwise. Moreover, this cutoff value increases in α: the higher it is, the more costly it is for A to lie. It first decreases in p_A because at $p_A \approx \frac{1}{2}$ because it becomes easier for A to afford lying completely as p_A rises. But eventually it may decrease in p_A (for low levels of α) or increase in it (for a sufficiently high α).

(1) To prove the first claim, fix signal quality p_A and α. If $\theta_A \leq \theta_A^2$, there exists a mixed strategy equilibrium. From IC (4), we know that x^d is the solution to $h(x^d, \theta_A) = 0$, where

$$\begin{aligned} h(x^d, \theta_A) &\equiv \frac{2p_A - 1}{1+(1-\theta_A)(1-x^d)} \\ &- \alpha\theta_A\Big[\frac{1}{\theta_A+(1-\theta_A)x^d} - \frac{p_A(1-p_A)}{1-p_A+p_A(1-\theta_A)(1-x^d)} - \frac{p_A(1-p_A)}{p_A+(1-p_A)(1-\theta_A)(1-x^d)}\Big] \end{aligned}$$

By the implicit function theorem, $\frac{dx^d}{d\theta_A} = -\frac{\partial h}{\partial \theta_A}/\frac{\partial h}{\partial x^d}$. From the proof of Proposition 1, $\frac{\partial h}{\partial x^d} > 0$, and $\frac{\partial h}{\partial \theta_A} < 0$ at $\theta_A \approx 0$, which implies that x^d increases in θ_A when θ_A is sufficiently small. Intuitively, if $\theta_A \approx 0$, a marginal increase in x^d improves A's reputation significantly, thus A becomes more honest. Next, suppose that $\theta_A \in (\theta_A^2 - \epsilon, \theta_A^2)$, then we can show that $\frac{\partial h}{\partial \theta_A} > 0$. Moreover, because $\frac{\partial^2 h}{\partial {\theta_A}^2} > 0$, the mixing probability x^d first increases in θ_A and then decreases in θ_A. Thus there exists a value θ_A^1 such that biased A is most honest if his prior objectivity $\theta_A = \theta_A^1$.

(2) To prove the second claim, fix prior objectivity θ_A and weight α. Similar to part (1), note that if $\alpha\theta_A < \frac{1}{2}$, then for some value of p_A, the LHS of IC (4) is always larger than the RHS even at $x^d = 0$. That is, if $p_A > \hat{p}_A$, A lies completely. The cutoff value $\hat{p}_A$ is implicitly defined such that $g(\hat{p}_A, \theta_A, \alpha) = 0$. If $\alpha\theta_A \geq \frac{1}{2}$, then we always have a mixed strategy equilibrium for all p_A. Next, IC (4) of biased A implicitly defines a function $f(x^d, p_A)$ such that x^d is the solution to $f(x^d, p_A) = 0$. We have:

$$\frac{\partial f}{\partial p_A} = \frac{2}{1 + (1-\theta_A)(1-x^d)} - \frac{(2p_A - 1)\alpha\theta_A(1-\theta_A)(1-x^d)[1 + (1-\theta_A)(1-x^d)]}{[1 - p_A + p_A(1-\theta_A)(1-x^d)]^2[p_A + (1-p_A)(1-\theta_A)(1-x^d)]^2}.$$

Clearly, if $p_A \approx \frac{1}{2}$, $\frac{\partial f}{\partial p_A} > 0$, thus $\frac{dx^d}{dp_A} < 0$. That is, if biased A's signal is very uninformative, he reports less truthfully as his signal quality improves. Simple algebra also show that $\frac{\partial^2 f}{\partial p_A^2} < 0$, thus $\frac{\partial f}{\partial p_A}$ decreases in p_A. How does A's reporting accuracy depends on p_A if the signal is highly informative depends on the value of α, θ_A. First, if $\alpha\theta_A < \frac{1}{2}$, then we can show that at the cutoff value $\hat{p}_A$, $\frac{\partial f}{\partial p_A} > 0$. In this case, the mixing probability x^d first decreases in p_A and becomes zero if $p_A > \hat{p}_A$. Second, if $\alpha\theta_A \geq \frac{1}{2}$, we can show that at $p_A \approx 1$, $\frac{dx^d}{dp_A} > 0$. Thus there exists a threshold quality $\overline{p}_A$ such that the equilibrium mixing probability x^d decreases with p_A for $p_A \in (\frac{1}{2}, \overline{p}_A]$ but increases if $p_A \geq \overline{p}_A$. $\square$

Proof of Corollary 2:

Suppose that $V_A(\pi_A) = \pi_A^\rho$ and $x^d > 0$ in the direct communication game where $\rho = 1$. Then similar to IC (4), the incentive constraint for the biased A when $s_A = 0$ becomes $g(x, \rho) = 0$, where

$$\begin{aligned} g(x,\rho) \quad \equiv \quad & \frac{2p_A - 1}{1 + (1-\theta_A)(1-x)} \\ - \quad & \alpha\Big[[\frac{\theta_A}{\theta_A + (1-\theta_A)x}]^\rho - p_A[\frac{(1-p_A)\theta_A}{1 - p_A(\theta_A + (1-\theta_A)x)}]^\rho - (1-p_A)[\frac{p_A\theta_A}{1 - (1-p_A)(\theta_A + (1-\theta_A)}]^\rho\Big]. \end{aligned}$$

If $\rho \approx 1, x \approx x^d$. Moreover, $\frac{dx}{d\rho} = -\frac{\partial g}{\partial \rho}/\frac{\partial g}{\partial x}$. At $\theta_A \approx 0, x^d \approx 0$, thus

$$\frac{\partial g}{\partial \rho} = -\alpha[ln(\frac{\theta_A}{\theta_A + (1-\theta_A)x})(\frac{\theta_A}{\theta_A + (1-\theta_A)x})^\rho + \infty] < 0.$$

Hence $\frac{dx}{d\rho} > 0$ if $\theta_A \approx 0$. Intuitively, A's posterior reputation after reporting $m_A = 1$ falls $\rho > 1$, because the convexity makes extremely low posteriors indistinguishable from zero. Hence A's net reputation cost actually increases and he needs to report more honestly.

Similarly, if $\theta_A \in (\hat{\theta}_A(p_A, \alpha - \epsilon, \hat{\theta}_A(p_A, \alpha)), x^d \approx 0$. Then it can be shown that $\frac{\partial g}{\partial \rho} > 0$, thus $\frac{dx}{d\rho} < 0$. Intuitively, reporting $m_A = 0$ is almost a sure sign of objective agent $Pr(A = o|m_A = 0) \approx 1$, thus the reputation changes little if $\rho \approx 1$. However, sending $m_A = 1$ is a risky gamble for biased A, thus his expected reputation increases if $\rho \approx 1$. This implies that his net reputation cost falls, and he lies more. □

Proof of Proposition 2:

Step 1: B's truth-telling incentive constraints. To find the equilibrium of the indirect communication game, first consider biased B's incentives after hearing m_A. On the one hand, B is concerned about C's action given his message. Let $a_1^B \equiv Pr(\eta = 1|m_B = 1), a_0^B \equiv Pr(\eta = 1|m_B = 0)$. Given the strategies of A and B described in the text, then $a_1^B - a_0^B$ is the marginal benefit B gets for reporting $m_B = 1$ versus $m_B = 0$:

$$a_1^B - a_0^B = \frac{2p_A - 1}{2 - (\theta_A + (1-\theta_A)x)(\theta_B + (1-\theta_B)y)}.$$

This net benefit increases in x and y, because the more honest A or B is, the more C believes in m_B.

On the other hand, agent B is concerned about how objective C thinks about him given his message and the true state. Specifically, B is concerned about how m_B affects his expected posterior reputation $E_\eta[Pr(B = o|m_B, \eta)|m_A] = \sum_\eta Pr(\eta|m_A)Pr(B = o|m_B, \eta)$. In particular, B's posterior objectivities given his message m_B and the (later) observed true state η are respectively:

$$\begin{aligned}
\tau_1 &\equiv Pr(B = o|m_B = 1, \eta = 0) = \frac{\theta_B[1 - p_A(\theta_A + (1-\theta_A)x)]}{1 - p_A(\theta_A + (1-\theta_A)x)(\theta_B + (1-\theta_B)y)}; \\
\tau_2 &\equiv Pr(B = o|m_B = 0, \eta = 0) = \frac{\theta_B}{\theta_B + (1-\theta_B)y}; \\
\tau_3 &\equiv Pr(B = o|m_B = 0, \eta = 1) = \frac{\theta_B}{\theta_B + (1-\theta_B)y}; \\
\tau_4 &\equiv Pr(B = o|m_B = 1, \eta = 1) = \frac{\theta_B[1 - (1-p_A)(\theta_A + (1-\theta_A)x)]}{1 - (1-p_A)(\theta_A + (1-\theta_A)x)(\theta_B + (1-\theta_B)y)}.
\end{aligned}$$

Combining these, we can show that, if $m_A = 0$, the net difference in B's posterior objectivity if he reports $m_B = 1$ as opposed to $m_B = 0$ is: $\tau_2 - p_A\tau_1 - (1-p_A)\tau_4$, which is positive, increasing in x but decreasing in y.

For biased B to report truthfully after $m_A = 0$ and $m_A = 1$ respectively, the following two incentive constraints must hold at $y = 1$:

$$a_1^B - a_0^B \leq \Delta_1 \equiv \beta \cdot [\tau_2 - p_A\tau_1 - (1-p_A)\tau_4]; \qquad (IC_1^B)$$

$$a_1^B - a_0^B \quad \geq \quad \Delta_2 \equiv \beta \cdot [\tau_2 - (1-p_A)\tau_1 - p_A \tau_4]. \qquad (IC_2^B)$$

Next, note that $\tau_2 > \tau_1$, $\tau_2 > \tau_4$ and $\tau_4 > \tau_1$, therefore the difference between biased B's *net* reputation cost is: $\Delta_1 - \Delta_2 = \beta(2p_A - 1)(\tau_4 - \tau_1) \geq 0$. This inequality shows that B's net reputation cost (of lying) after hearing $m_A = 1$ is always higher than that after hearing $m_A = 0$. Because even though a message of $m_B = 1$ is associated with bias, it is much worse for B's reputation if it turns out wrong. Moreover, observe from the incentive constraint IC_1^B above, the RHS is 0 while the LHS is strictly positive at $y = 1$, thus biased B never reports completely truthfully. It also implies that IC_1^B never holds strictly. For agent B, there can only be two possibilities: (IC_1^B) binds and (IC_2^B) holds strictly, in which case B mixes with probability $y > 0$ if $m_A = 0$ and report $m_B = m_A$ if $m_A = 1$; or (IC_1^B) does not hold, in which case B always lies.

Step 2: A's truth-telling incentive constraints. Agent A needs to compare his expected payoff after sending $m_A = 1$ versus $m_A = 0$, given B's strategy. Recall that x is the probability that he reports $s_A = 0$ truthfully. Then if $s_A = 0$, the net difference in A's expected payoff is:

$$\begin{aligned}
& EU_A(m_A = 1, s_A = 0) - EU_A(m_A = 0, s_A = 0) \\
= \quad & a_1^B + p_A Pr(A = o|m_B = 1, \eta = 0) + (1-p_A) Pr(A = o|m_B = 1, \eta = 1) \\
- \quad & Pr(m_B = 1|m_A = 0)[a_1^B + p_A Pr(A = o|m_B = 1, \eta = 0) + (1-p_A) Pr(A = o|m_B = 1, \eta = 1)] \\
- \quad & Pr(m_B = 0|m_A = 0)[a_0^B + Pr(A = o|m_B = 0)] \\
= \quad & Pr(m_B = 0|m_A = 0)[a_1^B + p_A Pr(A = o|m_B = 1, \eta = 0) + (1-p_A) Pr(A = o|m_B = 1, \eta = 1)] \\
- \quad & Pr(m_B = 0|m_A = 0)[a_0^B + Pr(A = o|m_B = 0)]
\end{aligned}$$

Observe first that both the net benefit and the net reputation cost is multiplied by a common factor: $Pr(m_B = 0|m_A = 0) = \theta_B + (1-\theta_B)y$, the probability that message $m_A = 0$ reaches C. The reason is that biased B may distort A's message, so both the improvement in A's objectivity and the loss in agenda pushing is affected similarly. Taking out the common factor from A's expected payoff, then A derives the same relative benefit from agenda pushing $a_1^B - a_0^B$. If $s_A = 0$, then A's relative reputation cost after reporting $m_A = 1$ is:

$$\alpha[Pr(A = o|m_B = 0) - p_A Pr(A = o|m_B = 1, \eta = 0) - (1-p_A) Pr(A = o|m_B = 1, \eta = 1)].$$

Moreover, A's posterior objectivities are respectively:

$$Pr(A=o|m_B=0,\eta=1)=Pr(A=o|m_B=0,\eta=0)=\frac{\theta_A}{\theta_A+(1-\theta_A)x};$$
$$Pr(A=o|m_B=1,\eta=0)=\frac{\theta_A[1-p_A(\theta_B+(1-\theta_B)y)]}{1-p_A(\theta_A+(1-\theta_A)x)(\theta_B+(1-\theta_B)y)};$$
$$Pr(A=o|m_B=1,\eta=1)=\frac{\theta_A[1-(1-p_A)(\theta_B+(1-\theta_B)y)]}{1-(1-p_A)(\theta_A+(1-\theta_A)x)(\theta_B+(1-\theta_B)y)}.$$

A faces two incentive constraints. Arguments similar to those about agent B can be used to show that there are only two possibilities: either A always lies; or A reports $m_A=s_A$ if $s_A=1$, but reports $s_A=0$ truthfully only with probability x.

Step 3: equilibrium. We now characterize the equilibrium of the indirect communication game. To simplify notations, define the following functions of x,y:

$$\begin{aligned}\xi(x,y) &\equiv \frac{2p_A-1}{2-(\theta_A+(1-\theta_A)x)(\theta_B+(1-\theta_B)y)}\\ &- \alpha\Big[\frac{\theta_A}{\theta_A+(1-\theta_A)x}-\frac{p_A\theta_A[1-p_A(\theta_B+(1-\theta_B)y)]}{1-p_A(\theta_A+(1-\theta_A)x)(\theta_B+(1-\theta_B)y)}\\ &- \frac{(1-p_A)\theta_A[1-(1-p_A)(\theta_B+(1-\theta_B)y)]}{1-(1-p_A)(\theta_A+(1-\theta_A)x)(\theta_B+(1-\theta_B)y)}\Big]. \end{aligned} \tag{9}$$

$$\begin{aligned}\psi(x,y) &\equiv \frac{2p_A-1}{2-(\theta_A+(1-\theta_A)x)(\theta_B+(1-\theta_B)y)}\\ &- \beta\Big[\frac{\theta_B}{\theta_B+(1-\theta_B)y}-\frac{p_A\theta_B[1-p_A(\theta_A+(1-\theta_A)x)]}{1-p_A(\theta_A+(1-\theta_A)x)(\theta_B+(1-\theta_B)y)}\\ &- \frac{(1-p_A)\theta_B[1-(1-p_A)(\theta_A+(1-\theta_A)x)]}{1-(1-p_A)(\theta_A+(1-\theta_A)x)(\theta_B+(1-\theta_B)y)}\Big]; \end{aligned} \tag{10}$$

The truth-telling incentive constraints of A and B when $s_A=0$ and $m_A=0$ can then be rewritten into: $\xi(1,y)\leq 0$ and $\psi(x,1)\leq 0$. From the analysis of A,B's incentive constraints above, biased A,B always report information supporting their agenda truthfully. If $s_A=0$ or $m_A=0$, there are three possible types of equilibria: (1) a fully mixed strategy equilibrium in which both agents report truthfully with positive probability: $x>0,y>0$. (2) A pure strategy equilibrium in which both A,B lie completely: $x=y=0$. (3) A hybrid equilibrium in which one agent always lies, and the other reports truthfully sometimes: $x=0,y>0$; or $x>0,y=0$. We consider these possible types of equilibria in turn.

First, suppose that $\xi(0,0)<0,\psi(0,0)<0$, which occurs if $\alpha>\tilde{\alpha}$ and $\beta>\tilde{\beta}$. The cutoff values are defined such that $\xi(0,0)=0, \psi(0,0)=0$ respectively at $\tilde{\alpha}, \tilde{\beta}$:

$$\tilde{\alpha} \equiv \frac{(2p_A-1)(1-p_A\theta_A\theta_B)(1-(1-p_A)\theta_A\theta_B)}{(1-\theta_A)(2-\theta_A\theta_B)(1-2p_A(1-p_A)\theta_A\theta_B)}; \qquad \tilde{\beta} \equiv \frac{(2p_A-1)(1-p_A\theta_A\theta_B)(1-(1-p_A)\theta_A\theta_B)}{(1-\theta_B)(2-\theta_A\theta_B)(1-2p_A(1-p_A)\theta_A\theta_B)}.$$

That is, if α, β are sufficiently high, even if one agent lies completely, the other agent still prefers to report truthfully sometimes. If there exists a mixed strategy equilibrium, $\xi(x,y)=0$ implicitly define the best response of agent A to B's truth telling: $x^{BR}(y)$, and $\psi(x,y)=0$ implicitly define B's best response to A's: $y^{BR}(x)$. We can see that both best response functions are continuous. Because $\xi(1,0)>0, \psi(0,1)>0$, and $\xi(x,y)$ increases in x and $\psi(x,y)$ increases in y, it must be that $\xi(x,0)=0, \psi(0,y)=0$ for some x, y. This implies that A's best response to y satisfies $x^{BR}(0)\in(0,1), x^{BR}(1)\in(0,1)$. Also, B's best response to x satisfies $y^{BR}(0)\in(0,1), y^{BR}(1)\in(0,1)$. Finally, because $x, y\in[0,1]$, by the intermediate value theorem, the two best response functions must intersect. Therefore there exists some x, y such that $\xi(x,y)=0, \psi(x,y)=0$. This establishes that if $\xi(0,0)<0, \psi(0,0)<0$, there exists a mixed strategy equilibrium.

Second, for uniqueness, let x, y be a mixed strategy equilibrium. Then, $\frac{dx^{BR}}{dy} = -\frac{\partial\xi(x,y)}{\partial y} / \frac{\partial\xi(x,y)}{\partial x}$, $\frac{dy^{BR}}{dx} = -\frac{\partial\psi(x,y)}{\partial x} / \frac{\psi(x,y)}{\partial y}$. From the analysis above, $\frac{\partial\xi(x,y)}{\partial x}>0$, $\frac{\partial\psi(x,y)}{\partial y}>0$. Moreover, it can be shown that $\frac{\partial\xi(x,y)}{\partial y}<0$, $\frac{\partial\psi(x,y)}{\partial x}<0$. Together, this means that $\frac{dx^{BR}}{dy}>0, \frac{dy^{BR}}{dx}>0$. Therefore the best response functions of A and B are strictly increasing. Straightforward calculations can show that $\frac{\partial\xi(x,y)}{\partial x}\cdot\frac{\partial\psi(x,y)}{\partial y} - \frac{\partial\xi(x,y)}{\partial y}\cdot\frac{\partial\psi(x,y)}{\partial x}>0$. This guarantees that whenever A and B's best responses intersect, A's best response function always has a higher slope than that of B's. This rules out multiple equilibria involving mixed strategies: thus there exists a unique fully mixed equilibrium if $\xi(0,0)<0, \psi(0,0)<0$.

Third, because A and B receive the same benefit from agenda pushing, it is straightforward to see that if $\alpha=\beta$, and that $\theta_A=\theta_B$, the RHS of the above equations are equal. Thus if the agents are symmetric, they lie with the same probability in the unique equilibrium of this game. Finally, consider the case that they care differently about their reputation, wlog, let $\alpha>\beta$. Then if in equilibrium, $x\leq y$, then A's net reputation cost is higher than that of B's, which is impossible. The only possibility is $x>y$. This shows that if $\theta_A=\theta_B$, the agent who cares about his reputation more reports more truthfully.

Next, if both α, β are sufficiently small, or if both θ_A and θ_B are sufficiently close to 1, then $\xi(0,0)\geq 0, \psi(0,0)\geq 0$, then this game has a pure strategy equilibrium in which the agents always report $m_A=1$, $m_B=1$. In this case, an agent prefers lying regardless of the other agent's report.

Finally, if α is sufficiently close to 0, but $\beta > \tilde{\beta}$, then $\xi(0,0) \geq 0, \psi(0,0) < 0$. Because $\psi(x,y)$ increases in y, and $\psi(0,1) \geq 0$, there exists a unique y such that $\psi(0,y) = 0$. In this case, A always reports $m_A = 1$ while B reports $m_B = m_A$ if $m_A = 0$ with probability y. Similarly, if $\xi(0,0) < 0, \psi(0,0) \geq 0$, then in equilibrium, $x > 0, y = 0$ such that $\xi(x,0) = 0, \psi(x,0) > 0$. □

Proof of Proposition 3:

Suppose that agent B becomes more concerned about his perceived objectivity. How does a small increase in β affect the equilibrium behavior of both agents? Recall that A, B's mixing constraints are given in Equation (9) and (10) respectively: $\xi(x,y) = 0$ and $\psi(x,y;\beta) = 0$. Let ξ_1 be the partial derivative of f with respect to its first argument x, and so on. Differentiate with respect to β, then:

$$\xi_1 x' + \xi_2 y' = 0, \quad \psi_1 x' + \psi_2 y' + \psi_3 = 0,$$

Solving these, the mixing probabilities change with a change in β in the following way:

$$\begin{cases} \frac{dx}{d\beta} = \frac{\psi_3 \xi_2}{\xi_1 \psi_2 - \xi_2 \psi_1}; & \text{indirect effect on } A \\ \frac{dy}{d\beta} = -\frac{\psi_3 \xi_1}{\xi_1 \psi_2 - \xi_2 \psi_1}, & \text{direct effect on } B. \end{cases}$$

Signs of some of the above partial derivatives are straightforward, namely, $\xi_1 > 0, \psi_2 > 0, \psi_3 < 0$. From proof of Proposition 2 above, we know also that in a mixed strategy equilibrium, $\xi_2 < 0$ and $\psi_1 < 0$. That is, even though both the benefit of agenda pushing and one's reputation cost increase when the other agent becomes more honest, the net effect is strictly negative. Also, we know that $\xi_1 \psi_2 - \xi_2 \psi_1 > 0$. This shows that the product of each agent's own response to changes in his honesty is larger than the product of his response to the other agent's changes in honesty. Therefore both x, y increases in β if there exists a fully mixed strategy equilibrium.

In addition, if $\theta_A \approx \theta_B$, it can be shown that $\xi_1 > |\xi_2|$, $|\psi_1| < \psi_2$. From expressions above, we can see that $0 < x' < y'$. If the agents have similar prior reputation, then even though both A and B's incentives to report truthfully increase in β, B responds more to the increase in his own reputational concerns than A does. The biased agents respond to a change in α similarly. □

Proof of Proposition 4:

First, consider the ex ante expected payoff of biased agent A. Suppose that there exists a mixed strategy equilibrium if he chooses to communicate directly, that is, $x^d > 0$, then we have:

$$\begin{aligned} EU_A^d &= \frac{1}{2}EU_A^d(s_A = 0) + \frac{1}{2}EU_A^d(s_A = 1) \\ &= Pr(\eta = 1|m_A = 1) + \frac{\alpha}{2}\Big[Pr(A = o|m_A = 1, \eta = 1) + Pr(A = o|m_A = 1, \eta = 0)\Big] \quad (11) \end{aligned}$$

The second equality is true because in equilibrium, biased A is indifferent between reporting $m_A = 1$ or $m_A = 0$ if $s_A = 0$, but he always reports $m_A = 1$ if $s_A = 1$.

Second, recall that a biased agent's payoff is equivalent to a weighted sum of two posteriors beliefs of C. We can show that:

$$\begin{aligned} & EU_A^d - (\frac{1}{2} + \alpha\theta_A) \\ = \;& Pr(\eta = 1|m_A = 1) - \frac{1}{2} + \frac{1}{2}\alpha\Big[Pr(A = o|m_A = 1, \eta = 1) + Pr(A = o|m_A = 1, \eta = 0) - 2\theta_A\Big] \\ = \;& \frac{(2p_A - 1)(\theta_A + (1-\theta_A)x^d)}{2[2 - (\theta_A + (1-\theta_A)x^d]} - \frac{\alpha\theta_A(1-\theta_A)(1-x^d)[1 - 2p_A(1-p_A)(\theta_A + (1-\theta_A)x^d)]}{2[1 - p_A(\theta_A + (1-\theta_A)x^d)][1 - (1-p_A)(\theta_A + (1-\theta_A)x^d)]} \\ = \;& 0. \end{aligned}$$

The last equality is simply the mixing condition (8) after rearranging terms. Therefore whenever biased A is mixing in direct communication, his ex ante expected payoff is simply the weighted sum of C's prior beliefs, by the law of iterated expectations. Similarly, whenever $x > 0$, biased A's ex ante expected payoff from indirect communication:

$$EU_A^i = Pr(\eta = 1|m_B = 1) + \frac{\alpha}{2}[Pr(A = o|m_B = 1, \eta = 1) + Pr(A = o|m_B = 1, \eta = 0)], \quad (12)$$

is also equal to $\frac{1}{2} + \alpha\theta_A$. Therefore if $x^d > 0, x > 0$, biased A is indifferent in ex ante terms between these two channels.

Third, if biased A strictly prefers lying, e.g. $x^d = 0$, then his ex ante expected payoff is strictly higher than the prior: $EU_A^d > \frac{1}{2} + \alpha\theta_A$. In this case, A's reputational cost is so low that $EU_A^d(m_A = 1|s_A = 0) > EU_A^d(m_A = 0|s_A = 0)$, thus he is worse off if he reports truthfully with any infinitesimally small $x^d > 0$.

(4.1) Suppose that A prefers lying in both channels: $x^d = x = 0$. Recall from the proof of Proposition 1 and 2 that $x^d = 0$ if $\alpha < \overline{\alpha}$ and $x = 0$ if $\alpha < \tilde{\alpha}$. Moreover, it is simple to show that $\overline{\alpha} < \tilde{\alpha}$. Thus for all $\alpha < \overline{\alpha}$, $x^d = x = 0$.

Next, given equations (11) and (12), we can show that, up to a factor $\theta_A(1-\theta_B)$:

$$\begin{aligned} & EU_A^d(s_A) - EU_A^i(s_A) \\ = \quad & \frac{2p_A - 1}{(2-\theta_A)(2-\theta_A\theta_B)} \\ - \quad & \frac{\alpha}{2}(1-\theta_A)(1-y)\Big[\frac{p_A}{(1-p_A\theta_A)(1-p_A\theta_A\theta_B)} + \frac{1-p_A}{(1-(1-p_A)\theta_A)(1-(1-p_A)\theta_A\theta_B)}\Big]. \end{aligned} \tag{13}$$

Clearly, if $\alpha \approx 0$ or $\theta_A \approx 1$, the second half of Expression (13) approaches zero. That is, any difference in A's net reputation cost in these channels becomes negligible. The first half of Expression (13), however, shows that the agenda pushing effect is strictly positive. Therefore biased A is strictly better off using direct communication.

Also, from Expression (13), if intermediary B is so truthful that $y \approx 1$, A chooses direct communication because an extremely truthful intermediary saves A very little in term of reputation cost, which is smaller than the loss of agenda pushing effectiveness. This occurs if β is sufficiently high. If A's weight on his reputation $\alpha \approx \overline{\alpha}$, and $p_A \approx 1$, we can show that $EU_A^d - EU_A^i$ is strictly negative if y is sufficiently small. This occurs if β is very low. biased A chooses indirect communication because if C attributes most of the blame to B if she receives a wrong message $m_B = 1$ but $\eta = 0$.

(4.2) Suppose that $x^d > 0, x = 0$, from the discussion above, indirect communication makes biased A better off because the reputation cost of lying completely is still small enough for him to lie completely ($x = 0$). But he only receives the prior if he communicates directly. This occurs, for instance, if $\overline{\alpha} \leq \alpha < \tilde{\alpha}$.

(4.3) Here the biased A receives the same ex ante expected payoff. However, if $s_A = 0$, biased A is indifferent between reporting $m_A = 0$ or $m_A = 1$. Therefore A's expected payoff after reporting $m_A = 0$ directly is simply $1 - p_A + \frac{\alpha\theta_A}{\theta_A + (1-\theta_A)x^d}$, and $1 - p_A + \frac{\alpha\theta_A}{\theta_A + (1-\theta_A)x}$ if intermediary B is used. Corollary 3 shows that $x^d > x$, hence A receives a higher expected payoff using the intermediary if the signal does not support his agenda. Consequently, it must also be that $EU_A^d(s_A = 1) > EU_A^i(s_A = 1)$: if the signal supports his agenda, A is better off if he has committed to direct communication. □

Proof of Proposition 5:

First, consider the k agents model where, for simplicity, $p_1 = 1, \theta_i = \theta, \alpha_i = \alpha$ for all agents. Similar to the proof of Proposition 1 and 2, each biased agent i faces two truth-telling constraints ($s_1 = 0$ and

$s_1 = 1$ respectively for agent 1):

$$EU_i(m_i = 1, m_{i-1} = 0) \leq EU_i(m_i = 0, m_{i-1} = 0);$$
$$EU_i(m_i = 1, m_{i-1} = 1) \geq EU_i(m_i = 0, m_{i-1} = 1).$$

Suppose that each biased agent i reports $m_{i-1} = 1$ truthfully ($s_1 = 1$ for agent 1), but $m_{i-1} = 0$ truthfully only with probability x_i ($s_1 = 0$ for agent 1). Then, if agent i hears $m_{i-1} = 0$ ($s_1 = 0$ for agent 1), the difference in his expected utility if he reports $m_i = 1$ instead of $m_i = 0$ is:

$$\begin{aligned} & EU_i(m_i = 1, m_{i-1} = 0) - EU_i(m_i = 0, m_{i-1} = 0) \\ = \ & Pr(m_k = 0|m_i = 0)\Big[Pr(\eta = 1|m_k = 1) - Pr(\eta = 1|m_k = 0)\Big] \\ - \ & Pr(m_k = 0|m_i = 0)\Big[\alpha Pr(i = o|m_k = 0) - \alpha Pr(i = o|m_k = 1, \eta = 0)\Big]. \end{aligned}$$

We can see that all biased agents derive the same net benefit from agenda pushing relative to their net reputation costs. The net benefit is the change in the decisionmaker's action induced by agent k:

$$Pr(\eta = 1|m_k = 1) - Pr(\eta = 1|m_k = 0) = \frac{1}{2 - \prod_i(\theta + (1-\theta)x_i)}.$$

Let $j \neq i$, thus each agent j reports m_{j-1} truthfully with probability x_j. Then the net reputation cost for agent i is:

$$\begin{aligned} & \alpha[Pr(i = o|m_k = 0) - Pr(i = o|m_k = 1, \eta = 0)] \\ = \ & \alpha\Big[\frac{\theta}{\theta + (1-\theta)x_i} - \frac{\theta(1 - \prod_j(\theta + (1-\theta)x_j))}{1 - \prod_i(\theta + (1-\theta)x_i)}\Big] \\ = \ & \frac{\alpha\theta(1-\theta)(1-x_i)}{(\theta + (1-\theta)x_i)(1 - \prod_i(\theta + (1-\theta)x_i))}. \end{aligned}$$

Similar to Proposition 1 and 2, there are only two possibilities for agent i: either always reports $m_i = 1$, or reports $m_{i-1} = 1$ truthfully, but reports $m_{i-1} = 0$ truthfully with probability x_i.

If α is sufficiently high, or θ is sufficiently small such that $\alpha(1-\theta) > \frac{1-\theta^k}{2-\theta^k}$, there exists a mixed strategy equilibrium such that for each agent i:

$$\frac{1}{2 - \prod_i(\theta + (1-\theta)x_i)} = \frac{\alpha\theta(1-\theta)(1-x_i)}{(\theta + (1-\theta)x_i)(1 - \prod_i(\theta + (1-\theta)x_i))}.$$

Observe from this mixing condition that all agents report truthfully with the same probability, $x_i = x_k$ for all i, are clearly an equilibrium. Moreover, note that this equilibrium is unique. For any two agents l and $l+1 \leq k$, suppose that $x_l > x_{l+1}$, then agent l receives the same net benefit from reporting $m_l = 1$, but pays a smaller net reputation cost than agent $l+1$, thus they cannot both be mixing, which is a contradiction. Similarly, in a $k+1$ symmetric agents model, $x_i = x_{k+1} > 0$ if $\alpha(1-\theta) > \frac{1-\theta^{k+1}}{2-\theta^{k+1}}$.

Second, to compare x_k and x_{k+1}. Suppose that $x_k = x_{k+1}$, then for any agent $i \leq k$, the difference in the decisionmaker's action after receiving a positive message becomes:

$$Pr(\eta = 1|m_k = 1) - Pr(\eta = 1|m_{k+1} = 1) = \frac{(\theta + (1-\theta)x_k)^k (1-\theta)(1-x_k)}{(2-(\theta+(1-\theta)x_k)^k)(2-(\theta+(1-\theta)x_{k+1})^{k+1})}.$$

The difference in the same agent's net reputation cost becomes:

$$\frac{\alpha\theta(1-\theta)(1-x_k)(\theta+(1-\theta)x_k)^k(1-\theta)(1-x_k)}{(\theta+(1-\theta)x_k)(1-(\theta+(1-\theta)x_k)^k)(1-(\theta+(1-\theta)x_k)^{k+1})}.$$

Next, let EU_i^k, EU_i^{k+1} denote respectively agent i's expected utility when there are k and $k+1$ agents. Compare the differences in his expected utility and use his equilibrium mixing condition, we can show that at $x_k = x_{k+1}$:

$$\begin{aligned}
& EU_i^k(m_i = 1, m_{i-1} = 0) - EU_i^k(m_i = 0, m_{i-1} = 0) \\
- \ & [EU_i^{k+1}(m_i = 1, m_{i-1} = 0) - EU_i^{k+1}(m_i = 0, m_{i-1} = 0)] \\
= \ & Pr(\eta = 1|m_k = 1) - Pr(\eta = 1|m_{k+1} = 1) - \alpha[Pr(i = o|m_{k+1} = 1, \eta = 0) - Pr(i = o|m_k = 1, \eta = 0)] \\
< \ & 0.
\end{aligned}$$

Note that agent i is mixing in both cases, this difference should be zero, which is a contradiction. Because i's expected utility after hearing $m_{i-1} = 0$ strictly increases in x^k, the only possibility is for $x_k > x_{k+1}$. Intuitively, the decrease in i's influence on the decisionmaker in term of agenda pushing is strictly smaller than the reduction in reputation cost for him. Thus if biased i lies in both cases with some probability, he lies more when there are $k+1$ agents. □

References

Arrow, K. J. (1985): "Information Structure of the Firm," *Papers and Proceedings of the Ninety-Seventh Annual Meeting of the American Economic Association*, 75(2), 303–307.

Arrow, K. J., and R. Radner (1979): “Allocation of Resources in large Teams,” *Econometrica*, 47, 361–385.

Bénabou, R., and G. Laroque (1992): “Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility,” *Quarterly Journal of Economics*, 107, 841–77.

Blume, A., O. Board, and K. Kawamura (2007): “Noisy Talk,” University of Pittsburgh, mimeo.

Chevalier, J., and G. Ellison (1999): “Career Concerns of Mutual Fund Managers,” *Quarterly Journal of Economics*, 114(2), 389–432.

CNN (2000): “‘Push polling’ takes center stage in Bush-McCain South Carolina fight; Dems campaign in California,” Cable News Network, February 10, 2000.

——— (2006a): “Libby court papers: Cheney said Bush OK’d intelligence leak,” Available at: http://www.cnn.com/2006/POLITICS/04/06/libby.ap/index.html.

——— (2006b): “Spector: White House needs to explain leak,” Available at: http://www.cnn.com/2006/POLITICS/04/09/whitehouse.leak/index.html.

Crawford, V. P. (2003): “Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions,” *American Economic Review*, 93(1), 133–149.

Crawford, V. P., and J. Sobel (1982): “Strategic Information Transmission,” *Econometrica*, 50(6), 1431–1451.

Dekel, E., and M. Piccione (2000): “Sequential Voting Procedures in Symmetric Binary Elections,” *Journal of Political Economy*, 108(1), 34–55.

DeMarzo, P., D. Vayanos, and J. Zwiebel (2003): “Persuasion Bias, Social Influence, and Unidimensional Opinions,” *Quarterly Journal of Economics*, 118(3), 909–968.

Dewatripont, M., and J. Tirole (1999): “Advocates,” *Journal of Political Economy*, 107, 1–39.

Evans, J. H., R. L. Hannan, R. Krishnan, and D. V. Moser (2001): “Honesty in Managerial Reporting,” *The Accounting Review*, 76(4), 537–559.

Goltsman, M., J. Horner, G. Pavlov, and F. Squintani (2007): "Mediation, Arbitraion and Negotiation," University of Western Ontario, mimeo.

Hendricks, K., and R. P. McAfee (2006): "Feints," *Journal of Economics and Management Strategy*, 15(2), 431–456.

Lewis, A. N. (2007): "Libby Guilty of Lying in C.I.A. Leak Case," *New York Times*.

Li, H., S. Rosen, and W. Suen (2001): "Conflicts and Common Interests in Committees," *American Economic Review*, 91(5), 1478–1497.

Li, W. (2006): "Peddling Influence through Well Informed Intermediaries," mimeo, UC Riverside.

——— (2007): "Changing One's Mind when the Facts Change: Incentives of Experts and the Design of Reporting Protocols," *Review of Economic Studies*, 74(4), 1175–1194.

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

Mirman, L. J., L. Samuelson, and E. E. Schlee (1994): "Strategic Information Manipulation in Duopolies," *Journal of Economic Theory*, 62(2), 363–384.

Morris, S. (2001): "Political Correctness," *Journal of Political Economy*, 109(2), 231–265.

Okrent, D. (2005): "Briefers and Leakers and the Newspapers Who Enable Them," The New York Times, May 8 2005.

Ottaviani, M., and P. Sorensen (2006): "Professional Advice," *Journal of Economic Theory*, 126(1), 120–142.

Prat, A. (2005): "The Wrong Kind of Transparency," *American Economic Review*, 95(3), 862–877.

Prendergast, C., and L. Stole (1996): "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, 104(6), 1105–1134.

Scharfstein, D. S., and J. C. Stein (1990): "Herd Behavior and Investment," *American Economic Review*, 80(3), 465–479.

Shanker, T., and E. Schmitt (2004): "Pentagon Weighs Use of Deception in a Broad Area," The New York Times, December 13, 2004.

Shapiro, C. (1986): "Exchange of Cost Information in Oligopoly," *Review of Economic Studies*, 53(3), 433–446.