USC FBE APPLIED ECONOMICS WORKSHOP
presented by Wei Li
FRIDAY, May 2, 2008
1:30 pm - 3:00 pm, **Room: HOH-706**

Peddling Influence through Well Informed Intermediaries*

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

September, 2007

Abstract

A sender with private information often tries to influence the decisionmaker through well-informed intermediaries such as experts or critics. Both the sender and the intermediary may be independently objective or biased: the objective type passes on the most accurate information, while the biased type wants to push a particular agenda but also to appear objective. Although using one's own information is a sign of objectivity, the biased intermediary selectively incorporates the sender's information to push his agenda. The intermediary's truth-telling incentives always decrease in those of the senders. Thus regulations and laws aimed at improving truth telling of the sender lead the intermediary to distort more, which may strictly lower the decisionmaker's payoff. In contrast, the sender and the intermediary's truth-telling incentives are strategic complements if they report simultaneously and independently.

JEL classification: C70, D82, M31, M38

Keywords: strategic communication through intermediaries, indirect marketing, disclosure laws, reputational concerns.

*I am grateful for the insightful comments from Glenn Ellison, Hao Li, R. Preston McAfee, Marco Ottaviani, Edward Schlee, Jean Tirole, and the seminar participants at UCR, Fuqua School of Business, Duke University, John Hopkins University, ASU, the 2007 Duke/Northwestern/Texas IO conference, and the 83th Western Economic Association Annual Conference.

1 Introduction

In business, politics and everyday life, a party with private information often tries to influence the final user of such information through well-informed intermediaries. In electronic commerce, online customer opinions or reviews have been shown to boost sales significantly.[1] Some businesses even pay experts or bloggers to promote their image or products while pretending to be independent reviewers (Bandler 2005, NYT 2006). In the medical and health industry, serious questions have been raised concerning pharmaceutical companies who promote their drugs through physicians and medical researchers (NYT 2002, Saul 2006, Carey 2006). These companies exert influence through funding research and giving perks such as paid consultancy positions; through providing physicians with free samples and the company's own information about their products ("detailing"); and through "ghostwriting" of journal articles where the purported academic authors have done little of the actual research.[2]

These influence activities have become more prevalent in the recent years. One estimate showed that approximately $19 billion is spent annually by drug companies for marketing to doctors.[3] In particular, the amount of money spent "detailing" physicians has increased from $3.0 to $4.8 billion from 1996 to 2000.[4] Nearly 75 percent of physicians in a national poll said the information they received from pharmaceutical representatives was "very" useful (15 percent) or "somewhat" useful (59 percent).[5] Many studies also indicate that these influence activities are effective in changing physicians' prescription behaviors, and consequently patients' welfare (Avorn, Chen, and Hartley 1982, Watkins,

[1] In a 2007 Customer Engagement Survey conducted by the global marketing research firm ACNielsen, online consumer opinions represented the third most trusted form of advertising, after word-of-mouth opinions and newspaper advertisements. Another market research firm, E-consultancy, asked online retailers about the effects of adding customer-generated reviews and ratings. Seventy-seven percent said site traffic increased, and 42 percent reported a rise in the amount of money spent.

[2] One recent example concerns an Annals of Internal Medicine article on Merck's "Advantage" trial of Vioxx, which omitted some trial participants' deaths. The article's first author Jeffrey Lisse, a rheumatologist at the University of Arizona, said that Merck actually wrote the report, and that "Merck designed the trial, paid for the trial, ran the trial.....Merck came to me after the study was completed." The New York Times, "Evidence in Vioxx Suits Shows Intervention by Merck Officials", April 24, 2005.

[3] For details, please see "Health industry practices that create conflicts of interest: a policy proposal for academic medical centers," published in the Journal of the American Medical Association (Brennan, Rothman, Blank, Blumenthal, Chimonas, Cohen, Golden, Kassirer, Kimball, Naughton, and Smelser 2006).

[4] See IMS Health Inc. and Competitive Media Reporting, as reported by Kaiser Family Foundation, 2001.

[5] More than half (55 percent) of a group of "high-prescribing" doctors surveyed by the industry data tracking group ImpactRx said that drug representatives served as their primary source of information about newly approved drugs. Only 26 percent of the doctors mentioned medical journals as their first information source. See "Getting Doctors to Say Yes to Drugs: The Cost and Quality Impact of Drug Company Marketing to Physicians" by the BlueCross BlueShield Association. See http://www.bcbs.com/betterknowledge/cost/getting-doctors-to-say-yes.html.

Moore, Harvey, Carthy, Robinson, and Brawn 2003).[6]

This paper presents a model of strategic communication through well-informed intermediaries and examines its effect on the final users of such information. In the medical industry, for example, a pharmaceutical company may promote its drugs through physicians, medical researchers, or through direct advertising. How are the physicians and researchers influenced by the pharmaceutical company, especially if their information disagree? How do the pharmaceutical company's truth-telling incentives interact with those of the intermediary's? What is the net impact of such indirect communication on the patients (or the broader medical community)? When does the company prefer direct advertising instead? This paper addresses these questions by investigating the strategic interactions of the sender and the intermediary. It also applies these insights to study the effectiveness of measures aiming at improving reporting accuracy, which has implications for professional ethical rules, disclosure and campaign finance laws.

In this model, a sender receives a private signal about the state of the world and sends a message to an intermediary who also has an independent, private signal. Because this paper focuses on a well-informed intermediary who has expertise such as physicians and researchers; or has experience in a market for credence goods, the intermediary's signal is assumed to be more accurate than that of the sender's.[7] The intermediary in turn sends a message to the decisionmaker, who takes an action based on what she hears.[8] Afterwards, the true state becomes observable, and the decisionmaker forms her opinion of how truthful the agents have been. The sender and the intermediary may be independently one of two types: objective or biased. An objective agent always passes on the most accurate information he has, while a biased agent has a private agenda: he wants the decisionmaker to act in a particular way. In addition, a biased agent is concerned about his reputation: he wants the decisionmaker to think

[6] Further examples include that of General Motors Corp., which found itself spending $52 million in 2001 for prescriptions that physicians wrote for Prilosec, a widely promoted and expensive drug for severe and persistent heartburn — even though a subsequent analysis found that 91 percent of the patients receiving prescriptions had no prior prescription related to heartburn and no prior diagnosis of the problem. A recent analysis of National Ambulatory Medical Care Survey data found that since 1991, "physicians are increasingly turning to expensive, broad-spectrum agents, even when there is little clinical rationale for their use" compared to simpler, generic antibiotics. By 1998-99, more than one in five adult prescriptions for broad-spectrum antibiotics and one in seven prescriptions for children were for conditions that were primarily viral, such as the common cold or acute bronchitis.

[7] In a companion piece, Li (2007b) considers the case when the intermediary has little information of his own and acts as a pure intermediary. The strategic interactions between agents as well as policy implications are very different due to the lack of information aggregation considerations. These differences are discussed in more details in Section 5.

[8] Throughout this paper, the sender and the intermediary are male and the decisionmaker is female.

that he is objective after she has observed the true state.

The objective sender always reports his signal truthfully, and the objective intermediary does so as well because his signal is more accurate than any message of the sender's even without distortion. But the biased agent faces a tradeoff between appearing objective and inducing the decisionmaker to take an action in favor of his agenda. The first result is that a biased intermediary may selectively incorporate the sender's message if his signal does not support his bias. Specifically, he is more likely to lie against his own, superior information if the sender's message supports his agenda than if it does not. Moreover, the less the intermediary is concerned about his reputation, the more likely he lies against his own information (and sometimes both their information) to push his agenda. This gives rise to a *bad information aggregation* effect because the intermediary's signal is the most useful to the decisionmaker.

It may appear that since the decisionmaker cannot observe the sender's message, the sender and the intermediary should share the blame if the intermediary's message turns out wrong, which makes it easier for each agent to lie more than he would have reporting alone. In this model, however, the intermediary's truth-telling incentives always *decrease* in those of the sender's: if the sender reports more truthfully, it becomes more attractive for the intermediary to lie. To see this, note that first, the bad information aggregation effect implies that if the sender is more truthful, his message becomes more credible, thus the intermediary is less likely to suffer a severe reputation loss by following the sender, which makes it cheaper for him to lie. Second, note that because the objective intermediary always reports his signal, independence from the sender's influence is the best sign of objectivity: the intermediary cannot blame any mistake on being mislead by the wrong source. This gives rise to a *substitution* effect: if the sender is more truthful, the intermediary's perceived objectivity is *less* responsive to his message. The decisionmaker assigns him less credit for a correct message because he may have followed the more accurate message from the sender, and more credit for a wrong message. These effects work together in reducing the intermediary's reputation cost and enabling him to lie more.

The decisionmaker may want to increase the reputational cost of the sender to lie, for instance by strengthening regulations over the pharmaceutical industry. Paradoxically, due to the bad information aggregation effect, this may make her *strictly* worse off if both agents have low levels of reputational concerns. Without such a policy, both the sender and the intermediary distort their information with high probabilities. In particular, the sender's message is not credible and unlikely to be followed by the intermediary. With such a policy, the sender is more truthful, leading the intermediary to lie more against his higher quality signal. This loss of the intermediary's high quality signal means that the decisionmaker may receive a message against both agents' true signals with a higher probability, and makes worse decisions. The decisionmaker should target the intermediary instead, perhaps by strengthening medical board review process and monitoring disclosure of industry ties of the physicians and researchers. The resulting gain in reporting accuracy from the intermediary, whose signal is more accurate, outweighs any indirect effect on the sender (even if it is negative). If the intermediary has high levels of reputational concerns, however, the substitution effect makes these policy measures relatively ineffective. Because the intermediary already reports very truthfully, and the substitution effect implies that the sender becomes less truthful, thus the net marginal impact of such policy approaches zero.

In comparison, independent (and simultaneous) communication from both agents is considered where the decisionmaker receives two messages before taking an action. At first glance, it may seem that one agent's message imposes some discipline on how much the other lies due to reputational concerns. With independent communication, however, each biased agent pays his own reputation cost because he is evaluated based on his message and the later observed true state. Therefore the presence of another message only affects the (expected) marginal impact of each agent on the decisionmaker's action, which is shown to decrease in the other agent's reporting accuracy. The biased agents' truth-telling probabilities are strategic complements under independent reporting. Intuitively, if one of the agents lies less, it increases the probability that the decisionmaker hears a message against a biased agent's agenda. Therefore the biased agent's lie is more likely to be contradicted and becomes less credible, because the decisionmaker is less likely to believe in his agenda given a conflicting message. Overall, distortion becomes less effective while one's reputation cost is unaffected, thus a biased agent lies less. This very complementarity suggests that the decisionmaker may want to encourage independent reporting, especially if both agents have high levels of reputational concerns.

Because indirect and independent (simultaneous) communication are widely used in different environments, a related question is which environment gives a biased sender a higher ex ante expected payoff. This paper shows that if the sender cares little about his reputation and the intermediary is considered very biased, he prefers sending his own message. Because here the intermediary's message is not credible and the sender does not want to waste his information: his independent message can push the decisionmaker's action further in favor of his bias. But if both of them have sufficiently high levels of reputational concerns, the sender prefers using an intermediary instead.

This differs from many existing papers analyzing how a sender influences the receiver directly by manipulating the information he sends (Crawford and Sobel 1982, Austen-Smith 1990, Dewatripont and Tirole 1999, Morris 2001, Ottaviani and Sorensen 2006, among others). Several recent papers consider the case where intermediaries may be affected by a biased sender. Durbin and Iyer (2006) consider the case where intermediaries (advisors) may be bribed by an uninformed and biased third party to support the third party's bias. All advisors are corruptible, but the good advisor is more costly because his interest is more closely aligned with the decisionmaker. Moreover, if the advisors have reputational concerns, a bribe may be necessary for the advisor to report his true signal if it happens to favor the biased third party. This paper focuses on the information transmission aspect: a source can influence an intermediary by altering his confidence in the signal, and thus the decisionmaker's action. But it can be extended to the case where cash bribes are also employed. In a similar vertical structure but with a different focus, Inderst and Ottaviani (2007) study the incentive problems faced by an intermediary (such as a sales agent) if he has to perform two tasks: one for the seller and the other for the buyer. Because of the inherent conflict of interest, they show how the firms design the intermediary's compensation scheme to balance his incentive to increase sales and to provide good advice to the buyer.

From a social network perspective, this paper is also related to DeMarzo, Vayanos, and Zwiebel (2003), who show that the influence of one's action on others depend not only on his information accuracy, but also on his position in a given social network. Their work takes the orthogonal approach from the present paper: they focus on richer sets of social networks, allowing each agent to report truthfully and assuming that the agent has "persuasion bias", namely they fail to account for possible repetitions in the information that reaches them. Instead, this paper focuses on a very simple way of indirect communication and considers strategic agents who make rational inference of any information they receive given the possible source(s) and the bias involved.

This paper proceeds as follows. Section 2 sets up the model and Section 3 analyzes the indirect communication game and considers the normative question of improving reporting accuracy. In comparison, Section 4 studies independent and simultaneous reporting by the agents. Section 5 discusses several main assumptions and Section 6 concludes. All the proofs are contained in the Appendix.

2 The Model

There are three agents: A, B and C. Agent C is the decisionmaker, whose optimal action depends on the state of the world $\eta \in \{0, 1\}$. Each state occurs with equal probability. Agent C chooses an action $a \in \Re$ to maximize her payoff, which is simply assumed to be the quadratic loss function $-(a-\eta)^2$. Her optimal action is thus equal to the probability she attaches to $\eta = 1$. In the medical industry example, the true state refers to the effectiveness of a drug: it may be "useless" (state 0) or "useful" (state 1). And the decisionmaker is the patient(s) who needs to decide how much to rely on this drug.

Although C has no information about the state, she has access to a partially informed source: agent B. Agent $i = A, B$ receives an informative signal $s_i \in S_i = \{0, 1\}$ about the true state:

$$Pr(s_i = 1|\eta = 1) = Pr(s_i = 0|\eta = 0) = p_i > \frac{1}{2}.$$

These signals are independent conditional on the state. As mentioned in the introduction, B's signal is assumed to be more informative than that of A's: $p_B > p_A$, reflecting the fact that B either has experience or expertise in evaluating a product. For example, a physician is better at figuring out whether a drug is useful for her patients. In these settings, the decisionmaker prefers hearing from an independent expert even at some cost of information aggregation, which is more reasonable if A's private information is of low quality. Thus it is assumed that A's signal is relatively uninformative.[9]

In the main model, after receiving signal s_A, A has an opportunity to send a private message $m_A \in M_A = \{0, 1\}$ to B. Agent B then sends a private message $m_B \in M_B = \{0, 1\}$ to C, given m_A and his own signal s_B.[10] Information is assumed to flow only in one direction, from A to B to C. An agent can only observe the message sent directly to him. Moreover, all messages are assumed to be observable but unverifiable, thus no transfers can be made based on the messages.

Agent i may be either objective (type o) or biased (type b). Each agent's type is independently drawn from $t_i = \{o, b\}$: $Pr(i = o) = \theta_i$, $Pr(i = b) = 1 - \theta_i$. Parameter θ_i, which captures agent i's existing reputation, is referred to as i's prior objectivity in this paper. An objective agent is assumed

[9] All the results in this paper hold if in equilibrium, biased B is willing to push his agenda if his own signal supports it (as characterized in Section 3). One sufficient condition for this to occur is that A's signal quality is below a cutoff value: $p_A \leq \overline{p}_A$, which always exists and is defined in the proof of Proposition 2.

[10] Agent B's message is best thought of as a simple, "yes-or-no" type of recommendation, which is the simplest way to illustrate the direction of biased B's distortions given his information. Allowing B to convey both A and his own information is discussed further in Section 5.

to be behavioral: he always passes on the information he believes to be the most accurate given what he knows. This can be justified either on the grounds of professional ethics or institutional goals.[11] Clearly, objective A reports his only (thus best) piece of information truthfully: $m_A = s_A$. Although objective B has two pieces of information m_A, s_B, his message choice is also very simple:

Observation 1 *Objective B always reports his own signal regardless of A's message: $m_B = s_B$.*

B's recommendation is simply his signal, which is a better source of information than that of A's, whether A's message confirms or contradicts it.[12]

A biased agent has an agenda: he always favors action $a = 1$ regardless of his information, but he also wants the decisionmaker to believe that he is objective due to reputational concerns. Let π_i be the decisionmaker's posterior estimate of agent i's objectivity after she observes the true state. Then the payoff functions of biased agent A and B are given respectively by:

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

The first half is C's action given what she hears from B.[13] Clearly, the closer C's action is to $a = 1$, the favorite agenda of a biased agent, the better off he is. The second half of the payoff function shows that a biased agent wants to appear objective, which is a reduced form for his reputational concerns, e.g., a movie critic may not exert much influence on potential viewers if he is perceived to be biased. Note that A and B may care about C's impression for different reasons: $\alpha, \beta \in [0, \infty)$ are the respective weights A and B place on C's (eventual) view of their objectivity. Alternatively, the ratios $\frac{1}{\alpha}, \frac{1}{\beta}$ reflect the extent, or the intensity, of A, B's bias. Thus the lower is α (β), the more is A (B) considered a "partisan" in that he cares more strongly about pushing his agenda.

In summary, the indirect communication game is illustrated in Figure 1.

In this game, biased agent i chooses message m_i given his information: $m_A : S_A \to \Delta(\{0,1\})$; $m_B : S_B \times M_A \to \Delta(\{0,1\})$. Given message m_B, C chooses action $a = Pr(\eta = 1|m_B)$. Later, she rationally

[11] For instance, Lahey Clinic, one of the major U.S. adult care hospitals writes: "Because good ethics begins with good medicine, the patient must receive accurate medical information and must understand it."

[12] More precisely, $Pr(\eta = s_B) > Pr(\eta = m_A)$ for any message agent A, objective or biased, may choose. It also implies that within the confines of this model, information aggregation is less important to the decisionmaker than hearing from objective B, a well-informed expert.

[13] Because the primary focus of this paper is on biased A and B, how C uses her updated beliefs of the agents is left unmodeled. In settings where C's optimal future decisions are very sensitive to the precision of her information, Li (2007a) shows that this updated belief itself may have value.

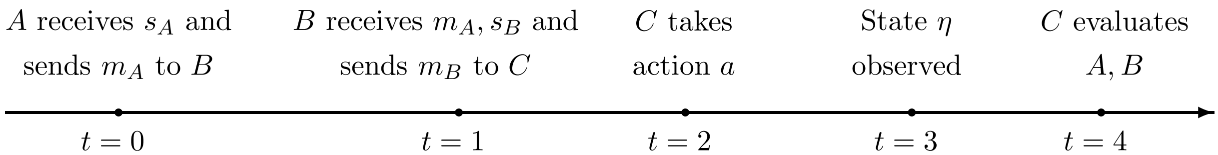

Figure 1: Timeline of the Indirect Communication Game

forms an opinion on A and B's posterior objectivity π_A, π_B as a function of their prior objectivity, the message received and the observed state. The equilibrium concept used in this paper is that of 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 in this model are private and unverifiable, this is not a cheap talk game and no babbling equilibria exist. The reason is that the objective agent always sends the most accurate information, hence any message is informative and useful to C, which implies that there is an endogenous cost of fabricating/passing on a biased message. When a biased agent pushes his agenda, he deviates from his best estimate of the state and thus is more likely to be wrong. In expectation, he is less likely to be considered objective in the eyes of C.

Before turning to the analysis, it is useful to keep in mind a few other interpretations of the model:

1. Agent A is a company launching a new movie, CD or some other consumer products. The state is the quality of the product. Agent B is a reviewer or critic who is better able to assess the quality. Agent C is the consumer who has no knowledge one way or another without the reviews.

2. Agent A is a mortgage broker (or banks and other lenders) and the state is the value of a property. If the broker (or the bank) is biased, it wants a higher appraisal so as to boost the commissions (or loans) from financing the purchase. Agent B is a real estate appraiser who specializes in estimating the market value of a property. The decisionmaker is the perspective homeowner who needs to decide how much, if any, to bid for the property.[14]

3. Agent A represents a grassroot political group who may be genuinely concerned about a policy or may be biased toward certain special interests. Agent B is a legislator overseeing this policy

[14] Biased appraisals have become a key issue in several recent lawsuits. For instance, New York's attorney general announced a case against eAppraiseIT, a leading appraisal management firm, for caving in to pressure from Washington Mutual to use a list of "proven appraisers" who he claims inflated home appraisals. Associated Press, November 1, 2007.

related area and makes a recommendation about this issue. Agent C is the voters (or its proxy the legislature) who need to know about the true consequences of this policy.

3 Peddling Influence through B

This section considers how communicating through a well-informed intermediary affects all biased agents' truthful reporting. Recall that an objective intermediary always reports his own signal because it is the best recommendation he could make. We now turn to the behavior of biased agents, who are partially trying to emulate the objective agents while pushing their agenda.

3.1 Biased Intermediary's Behavior

Although A has no influence on an objective intermediary, his message does have two effects on a biased one. The first is a bad information aggregation effect: because A's message contains information, it affects B's own estimates of the true state. These estimates in turn affect B's expected reputation, thus B's message may vary with A's — and become less useful for the decisionmaker. The second is on the credibility of B's message: if B may be influenced by A, the decisionmaker would take such influence into account before taking an action. This subsection analyzes how B's behavior depends on A's message and his own reputational concerns.

Suppose that biased A passes on a signal that supports his agenda truthfully, but one against his agenda probabilistically: $m_A = 1$ if $s_A = 1$, but $m_A = 0$ with probability $x \in [0, 1]$ if $s_A = 0$. Also, suppose that biased B chooses the following continuation strategy: reporting $m_B = 1$ if his signal supports his bias ($s_B = 1$). But if $s_B = 0$, he reports $m_B = 0$ with probability y_2 if A's message confirms his signal ($m_A = 0$); but reports $m_B = 0$ with probability y_1 if A's message contradicts his signal ($m_A = 1$). Both A and B's strategies are later shown to be part of the equilibrium strategies. Intuitively, if $y_1 \neq y_2$, the degree to which B is influenced by A depends on whether A's message concurs with his signal. The higher are the probabilities y_1, y_2, the more truthful B is. If $y_1 = y_2 = 1$, biased B reports his signal truthfully, like an objective agent; and if $y_1 = y_2 = 0$, he reports $m_B = 1$ to push his agenda regardless of what he knows.

Given these strategies, biased B chooses m_B to maximize:

$$EU_B(m_B|m_A, s_B) = Pr(\eta = 1|m_B) + \beta E_\eta[Pr(B = o|m_B, \eta)|m_A, s_B].$$

The first part is C's action given his message, and the second his expected posterior objectivity given his private information m_A and s_B, where the expectation is taken with respect to state η. To begin with, observe that the effectiveness of B's message is measured by the difference his message can induce in C's action: $Pr(\eta = 1|m_B = 1) - Pr(\eta = 1|m_B = 0)$. Because of the presence of objective B, this difference is always positive even if biased B lies completely. Also, A's message affects the effectiveness of m_B if B lets through A's information selectively. For example, $m_B = 0$ constitutes very strong evidence against biased agents' agenda: it indicates that both A and B's signals are likely to support $\eta = 0$.

Moreover, biased B is also concerned about how m_B affects his expected posterior objectivity:

$$E_\eta[Pr(B = o|m_B, \eta)|m_A, s_B] = \sum_\eta Pr(\eta|m_A, s_B)Pr(B = o|m_B, \eta).$$

Clearly, A's message influences B's estimate of the true state $Pr(\eta|s_B, m_A)$: it increases B's belief in his signal if they agree, but decreases it otherwise. The following proposition characterizes B's equilibrium behavior after receiving m_A.

Proposition 1 *Given x, there exists a unique continuation equilibrium in which biased B reports truthfully if his signal supports his agenda: $m_B = 1$ if $s_B = 1$. If B's signal is against his agenda ($s_B = 0$), the equilibrium is one of three possible types:*

(1) a total distortion equilibrium in which he always lies: $y_1 = y_2 = 0$; (2) or a strong distortion equilibrium in which he reports $m_B = 1$ if $m_A = 1$, and reports $m_B = 0$ sometimes if $m_A = 0$: $y_1 = 0, y_2 \in (0, 1]$; (3) or a weak distortion equilibrium in which he reports $m_B = 0$ if $m_A = 0$, and reports $m_B = 0$ sometimes if $m_A = 1$: $y_1 \in (0, 1), y_2 = 1$.

Proposition 1 shows that if $s_B = 0$, biased B turns to A's message for further evidence: he is more apt to lie if A's message supports, rather than contradicts, his agenda. Since m_A is always informative due to the presence of objective A, $m_A = 1$ lends some support for his agenda while $m_A = 0$ is another strike against it. This explains the different types of (continuation) equilibria: B lies completely if he has little reputational concerns. But if B is sufficiently concerned about his reputation, he is always more truthful if $m_A = 0$ than if $m_A = 1$: $y_2 > y_1$ if $\max\{y_1, y_2\} > 0$. To see this, note that for B, the benefit of agenda pushing is the same for all signals and messages, but it is more expensive in term of posterior objectivity to lie if $m_A = s_B = 0$ — lying against both their signals often leads to the worst reputation $Pr(B = o|m_B = 1, \eta = 0)$. If $m_A = 1$ instead, B believes less in his own signal, thus his expected reputation cost of lying is smaller. Consequently, if B reports $m_B = 1$ after $m_A = 0$ (even if just probabilistically), he must strictly prefer doing so after $m_A = 1$, which gives rise to a strong distortion equilibrium. Or, he may only report $m_B = 1$ sometimes if $m_A = 1$, which gives rise to a weak distortion equilibrium.

Which equilibrium may occur depends on B's reputational concerns, or equivalently, how strongly he feels about pushing his agenda. If β is sufficiently low ($\beta \leq \underline{\beta} \equiv \frac{1}{1+(1-\theta_B)}$), he always lies. As β increases, the continuation equilibrium is the strong distortion one: lying against both his highly informative signal and A's information is still not too costly in term of reputation. Eventually, as β becomes sufficiently high, the weak distortion equilibrium occurs. However, note that $y_1 < 1$: B never reports completely truthfully regardless of how high his reputational concerns are. If he were truthful, his message is credible, yet he pays no reputational cost for a wrong message, because C rightly believes in m_B and thinks that any mistake is due to a faulty signal, thus B strictly prefers lying to some extent.

Somewhat surprisingly, biased B's truth-telling incentives always *decrease* in x: the more truthful A is, the less truthfully B reports. We may have thought that since C cannot observe A's message, A and B should share the blame if the message turns out to be wrong. If so, B should lie less as well because if A is more truthful, C should attribute more blame to B, making it more costly for him to lie. In this model, however, the more truthful A is, the less likely m_A is biased toward $\eta = 1$, which has two effects. The first one is that $m_A = 1$ is more credible, therefore B is more likely to follow it and report $m_B = 1$ due to the bad information aggregation effect. Second, because biased A reports $m_A = 0$ with a higher probability and in equilibrium $y_2 > y_1$, it is more likely for biased B to report $m_B = 0$ as well if x increases. In this way, $m_B = 0$ becomes a less positive signal of independence (and hence B objectivity). Similarly, $m_B = 1$ becomes a less negative signal of B's objectivity because the chance of being influenced by $m_A = 1$ is smaller. Both these effects work together in reducing B's net reputation cost of agenda pushing, making it cheaper for B to lie if A is more truthful. At the same time, B's message becomes more credible because of A's more truthful reporting, hence B lies more.

Intuitively, any sign of being influenced is a sign of bias, which explains why B cannot shift any blame to A. As an example, note that C's posterior estimate $Pr(B = o|m_B = 0, \eta = 1) > Pr(B = o|m_B = 0, \eta = 0)$. That is, given $m_B = 0$, B is considered more objective even though his message is inaccurate. If $m_B = 0$ but $\eta = 1$, C thinks that it is likely that $m_A = 1$ — because it is more likely that $s_A = 1$ — but B did not follow it. This suggests that independent reporting may be a more prized sign of objectivity even though sometimes the prediction turns out wrong.

3.2 Equilibrium of the Indirect Communication Game

Having shown that biased B lets through A's message selectively, we now turn our attention to how biased A chooses a message m_A to maximize his expected payoff:

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

Biased A has a new consideration: he can only influence the decisionmaker C and later be judged for his objectivity through B's message. Despite this uncertainty over what the decisionmaker hears, the pivotal event for A, which determines his truthful reporting, is the difference in B's messages induced by m_A. Formally, the difference in A's expected payoff, $EU_A(m_A = 1|s_A = 0) - EU_A(m_A = 0|s_A = 0)$, is proportional to $(1 - \theta_B)(y_2 - y_1)$. Intuitively, if B is known to be objective ($\theta_B = 1$); or if B always lies to push his agenda ($y_1 = y_2 = 0$), A's message has no impact on either C's action, or his own reputation. The gap between y_2 and y_1 reflects the magnitude of A's influence. After all, A's message only matters if $s_B = 0$, in which case $m_A = 1$ is more likely to change B's message than $m_A = 0$ because $y_2 > y_1$. By the same token, in C's posterior estimates, A's net reputation cost depends more strongly on the event $\eta = 0$ than $\eta = 1$, because B is most likely to have received $s_B = 1$ if $\eta = 1$, in which case A has no influence.

Biased A and B compare their agenda pushing effectiveness against possible reputation losses. Their equilibrium behavior is given by the following:

Proposition 2 *Biased A reports truthfully if his signal supports his agenda: $m_A = 1$ if $s_A = 1$.*

(2.1) Biased A has no influence, $x \in [0, 1]$, if B plays a total distortion equilibrium, which occurs if B has sufficiently low reputational concerns, or if B's prior objectivity is sufficiently high.

(2.2) Biased A has influence, $x \in [0,1)$, if B plays a strong or weak distortion continuation equilibrium. A always reports $m_A = 1$ if α is sufficiently low or if θ_A is sufficiently close to 1. Moreover, $m_A = 1$ if β is sufficiently high.

Proposition 2 shows that A has no influence on the decisionmaker if biased B always reports $m_B = 1$. In this case, A's information is completely lost and he is free to report in any way: it does not affect the outcome in term of C's action, nor does it has any impact on his own reputation. Biased A only has influence if B may alter his message because of m_A, in which case A never reports completely truthfully because his information can affect C's action indirectly. More interestingly, if A's influence on B is sufficiently small, which occurs if β is sufficiently high, A always reports $m_A = 1$ regardless of his own reputational concerns. The reason is that biased B follows his own signal very closely due to his reputational concerns ($y_2 = 1, y_1 \approx 1$), thus C attributes very little blame to A regardless of m_B. In addition, B's message is highly credible because of his truthful reporting. Thus A lies completely: the indirect impact of his message on C's action strictly outweighs his rather negligible reputation cost.

In a weak distortion equilibrium, if A is sufficiently concerned about his reputation such that $x > 0$, A's truth-telling incentives decrease in y_1. To see this, first note that A's agenda pushing effectiveness increases in both agents' truthful reporting: $m_B = 1$ becomes more convincing if the agents are more honest. Second, due to the aforementioned substitution effect, A's net reputation cost decreases in y_1 because it becomes less sensitive to B's message, right or wrong. For instance, suppose that the true state $\eta = 0$, consider how A's posterior objectivity changes as a result of a higher y_1. On the one hand, $Pr(A = o|m_B = 0, \eta = 0)$ decreases in y_1 because C believes that $m_B = 0$ is more likely a sign that B is more honest and didn't follow $m_A = 1$, thus it is a less positive signal of A's objectivity. On the other hand, $Pr(A = o|m_B = 1, \eta = 0)$ increases in y_1, because if B is more truthful, C thinks that it is more likely that B receives the wrong signal instead of being influenced by A. As a result, the difference between these two posteriors–the major part of A's net reputation cost–decreases in y_1. Thus A lies more because he exerts a stronger influence on C at a lower reputation cost. Since B's truth-telling incentives always decrease in x, x and y_1 are strategic substitutes in a weak distortion equilibrium.

In a strong distortion equilibrium, however, A's response to B's truth telling depends on the parameter values because both A's agenda pushing effectiveness and his reputation cost increase in y_2. It is important to note that here biased B is (primarily) following m_A: if $m_A = 1$, $m_B = 1$; but if $m_A = 0$, and y_2 increases, biased B reports $m_B = 0$ more often. This means that A and B are (partially) sharing blame of message $m_B = 1$, because B most likely has simply followed A's message. This implies that A's reputation cost increases in B's truth telling. For instance, $m_B = 0$ is a good signal of A's objectivity because C believes that it is more likely that A sends $m_A = 0$ and that B has followed. Otherwise, B would have sent $m_B = 1$. Similarly, $Pr(A = o|m_B = 1, \eta = 0)$ decreases in y_2. This effect may be sufficiently strong that A reports more truthfully if B does so in a strong distortion equilibrium. Intuitively, this asymmetry in A and B's responses to the other's truth telling is driven by the fact that B can never shift any blame to A, while A may or may not be blamed for a wrong message of B's.

The equilibrium of this game shifts with B's reputational concerns. Recall that if $\beta \leq \underline{\beta}$, a total distortion equilibrium exists. When $s_B = 0$, there exists cutoff values β_1 and $\overline{\beta}$ such that if $\beta \in [\underline{\beta}, \beta_1]$, the strong distortion equilibrium occurs; and if $\beta \geq \overline{\beta}$, the weak distortion equilibrium occurs.[15] Interestingly, if biased B has moderate reputational concerns, e.g. $\beta \in [\beta_1, \overline{\beta}]$ and A's reputational concerns are not so low that he always reports $m_A = 1$, multiple equilibria involving both strong and weak distortions may exist. In this range, B's behavior is very sensitive to how accurate A's message is, especially if B's signal quality is close to that of A's. If A is very truthful, then $m_A = 1$ is relatively credible, which means that B is less likely to be wrong if he reports $m_B = 1$, giving rise to a strong distortion equilibrium. If, however, $m_A = 1$ is not very credible, then lying and reporting $m_B = 1$ is more expensive for B, thus a weak distortion equilibrium may exist as well.[16]

3.3 Information Loss for the Decisionmaker

If well-informed intermediaries are influenced by sources with inferior and potentially distorted information, the decisionmaker may want to discourage such influence activities to reduce her expected loss, which may be quite severe.[17] To encourage independent recommendations, the decisionmaker may make it more expensive, in term of reputation cost, for the intermediaries to lie by raising β; or she

[15] These cutoff values are defined in the proof of Proposition 2 in the Appendix.

[16] As a numerical example, if $\alpha = 10, \theta_A = \theta_B = 0.5, p_A = 0.7, p_B = 0.75$, then multiple equilibria exist for $\beta \in [1.55, 1.62]$.

[17] For instance, the settlement of the $185 million class action lawsuit against Bristol-Myers Squibb in January 2006 shows that they paid physicians to exaggerate, in major medical meetings, the benefits of their drug for patients with high blood pressure and heart failure; these physicians also failed to report publicly on substantial numbers of life-threatening drug complications which they knew, from their close relationship to the company, to exist.

may encourage the sources to be more truthful by raising α.[18] In real life, such measures may take the form of more stringent and thorough medical board reviews; or stricter regulations on pharmaceutical companies.

Recall that C's optimal action given m_B is to set $a = Pr(\eta = 1|m_B)$, thus her expected payoff is simply $-E_\eta E_{m_B}[(Pr(\eta = 1|m_B) - \eta)^2|m_B]$. Distortions by the biased agents affect the accuracy of m_B and thus her payoff in two ways: through the credibility of B's message; and through the overall probability a distorted message reaches her. The following result describes how changing the agent's reputational concerns affect their reporting accuracy through their strategic interactions.

Proposition 3 *There exists a cutoff value $\underline{\alpha}$ such that if $\alpha \geq \underline{\alpha}$, x increases in α while y_1 or y_2 decreases in it. In the weak distortion equilibrium, x decreases in β while y_1 increases in it. In the strong distortion equilibrium, if p_B is sufficiently high, x and y_2 increase in β.*

Proposition 3 suggests that measures increasing α may exacerbate the information loss of C if A and B both have low to moderate levels of reputational concerns. That is, increasing the reputation (or financial) cost of pharmaceutical companies influencing the physicians or medical researchers may make the decisionmaker *strictly* worse off. The reason is that even if A can influence B's message, he always reports $m_A = 1$ if his reputational concerns are sufficiently low: $\alpha \leq \underline{\alpha}$. In this case, his message is not very credible and the intermediary does not follow $m_A = 1$ closely in a strong distortion equilibrium. However, a small increase in α makes A more truthful, thus message $m_A = 1$ becomes more useful to biased B. The aforementioned bad information aggregation effect implies that B follows m_A more closely (y_2 falls in α). Because B's signal is more informative than A's, the ensuing information loss for the decisionmaker is more severe, and the net effect may be negative. For example, if $\theta_A = \theta_B = 0.5, p_A = 0.7, p_B = 0.95$, then for $\beta \in [0.7, 1]$, an increase from $\alpha = 1$ to $\alpha = 2$ reduces C's expected payoff. Intuitively, a small increase in the source's reputational concerns may worsen the intermediary's incentives when it counts the most. The decisionmaker should consider policies targeting the intermediary instead.

[18] Though we only consider measures affecting the reputational cost of the agents and ignore monetary fines or any costs required to implement and enforce such measures, such costs can be easily incorporated. For instance, if C's cost of imposing stronger regulations are sufficiently convex, then only small changes in α, β may be feasible, which may be counterproductive as implied by Proposition 3.

The fact that B's truthful reporting is more important than A's in a strong distortion equilibrium — because he may lie against his high quality signal — also explains why a small increase in β is more effective than α. When β increases, B lies less against his high quality signal. Moreover, if B's signal quality is sufficiently high, Proposition 3 shows that A lies less as well because C now attributes more blame to A when B's message turns out wrong, making it more expensive for A to lie. Clearly, if both are more truthful, C is more likely to receive an accurate report and her expected payoff increases. Specifically, any measure raising β yields the highest marginal benefit if $\beta \in \{\underline{\beta}, \beta_1\}$. In the medical industry example, strengthening medical board review process or disclosure rules may drastically improve the credibility of the medical profession: both because the physicians are less influenced and also because the drug companies become more truthful in revealing side effects.[19]

A small increase in β, however, has no effect on C if B has very low reputational concerns ($\beta \leq \underline{\beta}$), because B does not change his behavior in a total distortion equilibrium. Also, it only has a very small positive effect on C if B has very high reputational concerns ($\beta \gg \overline{\beta}$). Because in a weak distortion equilibrium, x and y_1 are substitutes, thus an increase in β leads A to lie more (or completely if β is sufficiently high). Moreover, in this region, because B already reports very truthfully, C's expected payoff increases very gradually in β; in fact, the marginal benefit of raising β approaches zero.[20] The small improvement in B's truthful reporting is partially offset by the loss if B's message, which is highly credible, is wrong: y_1 increases but the loss $(Pr(\eta = 1|m_B = 1) - 1)^2)$ also increases.

4 Comparing Communication Methods

For legal or institutional reasons, in environments such as marketing and lobbying, exerting influence through intermediaries is common; in other environments such as advertising, A conveys his information to the decisionmaker without intermediaries. Independent (and simultaneous) reporting serves as a natural benchmark against the indirect communication model above, because both agents have informative signals and their messages are potentially useful if they influence the decisionmaker directly. This section studies the biased agents' truth-telling incentives under independent reporting and how

[19] As discussed in Subsection 2, A may lie more as well in a strong distortion equilibrium if B becomes truthful. However, because B's signal is more precise, C still benefits from raising β even though the effect is less pronounced.

[20] Hence measures raising β for intermediaries sufficiently concerned about their reputations makes C worse off if a small fixed cost is involved.

these incentives affect the decisionmaker's expected payoff.

4.1 Independent (and Simultaneous) Reporting

Suppose instead of communicating through an intermediary, A sends a message to C, who now receives both m_A and m_B before taking an action. All the other assumptions remain. When biased A reports independently, he chooses m_A to maximize:

$$E_{m_B}[Pr(\eta = 1|m_A, m_B)|s_A] + E_\eta[Pr(A = o|m_A, \eta)|s_A].$$

Observe that the first part, A's influence on C, differs from the case of indirect communication because it now depends on B's message as well. But the second part, A's expected posterior objectivity, is independent of m_B. It may seem that the presence of another message may affect A by imposing additional discipline on his distortion, but since C has observed the true state when she evaluates an agent, A's posterior objectivity depends solely on his message. More precisely, the presence of multiple messages affects the tradeoff a biased agent faces primarily by changing the *marginal* impact of his message on the decisionmaker's action, and thus his agenda pushing effectiveness.

Under independent reporting, then, all the strategic interactions between the agents enter through their messages' influence on C's action; their posterior reputations are not affected like in the case of indirect communication. Each agent chooses a message after comparing the net marginal impact his message has on C's action with the net loss of reputation if he lies.

Proposition 4 *Biased agent $i = A, B$ reports $m_i = 1$ if $s_i = 1$. If $s_i = 0$, biased i always reports $m_i = 1$ if he attaches little weight to his reputation (α or β sufficiently low); or if his prior objectivity is very high (θ_A or θ_B sufficiently high). Biased i reports $m_i = 0$ truthfully with probability $x_i > 0$ otherwise. Moreover, the agents' truth-telling probabilities are strategic complements.*

The key of Proposition 4 is that, for each agent, the net benefit from agenda pushing decreases in the other agent's truth telling *when* it matters. Suppose that $s_A = 0$ and B reports more truthfully (x_B increases), it has two effects on A. If $m_B = 1$ as well, m_A's marginal impact on C's action decreases — the more truthful B is, the less responsive C is to A's message. However, if $m_B = 0$, which is more likely since $s_A = 0$, $m_A = 1$ has a stronger influence on C because she receives at least one message in support of his agenda. But it can be shown that the decrease in A's marginal impact in the event $m_B = 1$ dominates. Intuitively, the more C thinks that $\eta = 0$, the more message $m_A = 1$ appears to be distorted and less credible, and the less she changes her action. Because she is more likely to receive message $m_B = 0$ if $s_A = 0$, she is more likely to believe that $\eta = 0$, thus A's agenda pushing is less effective than if $m_B = 1$. In this way, A's (expected) marginal impact on C falls but his net reputation cost remains the same. Consequently, A lies less than he would have as the sole source of information. Moreover, this also implies that with independent, simultaneous reporting, A and B's truthful reporting probabilities x_A, x_B are strategic complements.[21]

In particular, this implies that in comparison with either indirect communication, or reporting as the sole source of information (only A or B sends a message), a biased agent is less likely to lie completely. Formally, biased i always reports $m_i = 1$ if he cares little about reputation, but the cutoff value is lower for him to start reporting truthfully than in the other two cases.[22] Proposition 4 also suggests that, if both A and B are sufficiently concerned about their reputations, measures increasing the cost for either agent to lie has a stronger positive effect on the agents' truth telling than that in the indirect case where the agents' truth-telling incentives are substitutes. Together, this suggests that policies aiming to improve reporting accuracy from potentially biased sources should take a two-pronged approach: increasing the cost of distortion by intermediaries if both A, B have low to moderate levels of reputational concerns; and encouraging independent reporting, especially if the agents have high levels of reputational concerns. One way of achieving this is to build some distance between the source and the intermediary, perhaps by reducing the direct linkage between medical researchers, doctors and individual pharmaceutical companies.[23]

One caveat concerning policies encouraging independent reporting is that it may be counterproductive if both agents have such low reputational concerns that they always push their agenda ($x_i = 0$). Because in this case, under indirect communication, C only receives one agenda-pushing message from

[21] This complementarity is driven by each agent's decreasing marginal impact on C instead of changes in reputation cost as in Li (2007b), where the uncertainty of who initiates a wrong message creates a blame sharing effect and encourages each agent to lie more.

[22] For instance, under independent reporting, biased B starts reporting truthfully with some probability if $\beta \geq \frac{p_A(1-p_A)(1-\theta_B)}{1-\theta_B+p_A(1-p_A)\theta_B^2}$, which is strictly smaller than the cutoff value $\underline{\beta}$ in the indirect case.

[23] One recent instance to insulate medical research and treatment from industry influence is when the University of Pennsylvania Health System announced that industry can make gifts to departments to support educational programs (but not to individual faculty), while the money is disbursed at the discretion of department chiefs and chairs.

biased B. Under independent reporting, however, she receives two agenda-pushing messages which reinforce and strengthen each other. Due to the possible presence of objective agents, C is more persuaded by these messages. If one of the agents is likely to be objective, this loss may be outweighed by the possibility of receiving one truthful message. But if both agents are very biased, this may exacerbate C's information loss. In situations where agents have a clear agenda and little reputational concerns, then, it may be better to discourage independent reporting by reducing the quantity of distorted information the decisionmaker receives, such as the amount of direct TV advertisements for prescription medicines.

4.2 Biased A's Preference

Having examined some implications of independent reporting on the decisionmaker's expected payoff, this subsection proceeds to consider a related question: the ex ante preference of a biased agent A. That is, whether independent reporting or communicating through B gives A a higher expected payoff if he chooses an environment before receiving his signal.[24] Knowing this helps us understand when peddling influence through experts is more useful to a biased source. For instance, when does a pharmaceutical company prefer direct advertising, and when is he better off communicating through physicians instead?

Biased A's ex ante expected utility if he communicates through B is simply:

$$E_{s_A}\Big[E_{m_B}[Pr(\eta = 1|m_B) + \alpha E_\eta[Pr(A = o|m_B, \eta)]]\Big|s_A\Big].$$

Note that if his message has no influence on B, then the first half of the above expression, C's (expected) action, purely depends on the characteristics of B; and the second half is just the prior, $\alpha\theta_A$, because his reputation is not affected at all. But if he can influence B's message, A needs to compare his impact on C's action with his loss in expected reputation.

Proposition 5 (Independent vs. Indirect Communication) *(5.1) If α and β are sufficiently low, A prefers independent reporting if θ_B is sufficiently close to* 0.

(5.2) If α and β are sufficiently high, A prefers using an intermediary.

Proposition 5 shows that if the intermediary, despite his better signal, is perceived to be very biased and not credible, A should try to persuade the decisionmaker directly instead of wasting his informative

[24] This is the case where biased A commits to his choice before observing his signal, thus his preference does not reveal information about his signal.

signal. The reason is that if A and B have sufficiently low reputational concerns, A's message has no influence on B, and in turn the decisionmaker, as shown in Proposition 2. Therefore A prefers an environment where C's action favors his agenda more. If B's message is not credible due to his low prior objectivity, A's message has a stronger impact on C's action than B's. Doing so leads to a lower posterior objectivity for A, which he cares little about.

Next, if in equilibrium, A reports truthfully with positive probabilities in both these environments and B does not lie completely, then the law of iterated expectation applies and his ex ante expected payoff is simply the sum of his priors: $\frac{1}{2} + \alpha\theta_A$. Hence he is indifferent. To see the second part of Proposition 5, recall that biased A always reports $m_A = 1$ if β is sufficiently high regardless of his own reputational concerns. In this case, B reports so truthfully ($y_1 \approx 1$) that A pays a negligible reputation cost (which is proportional to $1 - y_1$); but he still has some influence on the decisionmaker because B's message is highly credible. Thus he can "afford" to lie indirectly while he could not do so, at the same level of reputational concerns, if he sends a message of his own. Hence if B faces sufficiently high reputational concerns, A is better off influencing him than risking losing his own reputation. This suggests that a pharmaceutical company reasonably concerned with its reputation always prefers communicating through physicians or researchers strongly concerned about their own reputations.

5 Discussions

Having examined the incentive problems biased A and B face in different environments involving either indirect communication or independent reporting, this section discusses several main assumptions on the agents' information quality, preferences of the objective agent, as well as the distribution of state η. It also suggests how agents' behavior may change if these assumptions were varied.

A. Relative informativeness of signals. B's signal is assumed to be more accurate than that of A's, which implies that B is considered more objective if he appears free from A's influence. This model thus fits more closely the environments where the intermediary possesses good information such as experts or critics. In some environments, agent A has superior and/or exclusive information, for example, government policy or military intelligence. In that case, A's signal may be far more informative and an objective intermediary behaves differently. He follows A if it is unlikely that A has lied, but may dismiss it if biased A lies a lot: $m_B(m_A = 1, s_B = 0) = 0$ if $x \leq x_1$, $m_B(m_A = 1, s_B = 0) = 1$ otherwise. The cutoff value x_1 increases in p_B, and in particular, if B's signal is sufficiently uninformative, an objective B always follows m_A despite A's potential distortion. Li (2007b) shows that in this case, B serves as a pure intermediary, and biased A and B's truth telling become complements because they share the blame of any distortion. Moreover, agent A always lies more using an intermediary — and he may prefer a more biased intermediary to a more objective one — because the decrease in his reputation cost strictly outweighs any loss in his message's credibility due to B's possible distortion.

B. Message space of agent B. The well-informed intermediary here sends a simple message $m_B \in \{0, 1\}$ to the decisionmaker C. One question is why the intermediary does not indicate both A's message as well as his own signal. One reason for restricting intermediary's message space is that often experts only convey their recommendations as opposed to intensity to the decisionmaker, especially if the issue at stake is complex, for example whether to take a particular medicine. Another reason is that this restriction helps illustrate the direction of B's distortions: the bad information aggregation effect still exists and may even be intensified with a richer message space. The objective intermediary reports a vector of messages (m_A, s_B) truthfully, but the biased type still has incentives to push one or both messages toward 1.[25] Giving B a richer message space may not reduce the information loss of the decisionmaker. For instance, if B's reputational cost is sufficiently low, he may report $m_B = (1, 1)$ if $s_B = 0$, which is more effective than $m_B = 1$ because it may be true signals from two objective agents.

C. Role of the objective types. In this model, the objective agent has a preference for accuracy: he conveys the best information available to him. This assumption, consistent with the stated goals of experts such as stock analysts, critics and physicians, also greatly simplifies the inference problem of decisionmaker C. An objective agent, however, may be concerned about his reputation as well as passing on accurate information such that an objective agent i chooses m_i to maximize $Pr(\eta = s_i|m_i) + \lambda Pr(i = o|m_i, \eta)$, given his signal s_i. Morris (2001) shows in a model of direct communication that an objective expert may lie even though he wants the decisionmaker to take the correct action, because she doesn't want to be confused with a biased type and lose future influence. In the present model, similar incentives may arise if the objective intermediary is also concerned about his reputation. Recall that $m_B = 0$ but $\eta = 1$ is a good sign of B's objectivity because it suggests that he is not influenced by $m_A = 1$. Thus an

[25] Sixteen incentive constraints, instead of four, must be satisfied for biased B to report truthfully.

objective B has a stronger incentive to report 0 if $m_A = 0, s_B = 1$ than if he reports alone, which may drive A to report $m_A = 1$ more often, knowing that the objective B may avoid $m_B = 1$ on his own.

D. Asymmetric state distribution. In this model, the states are assumed to be ex ante symmetric, but the agents may believe that one state is more likely than the other: $Pr(\eta = 0) > \frac{1}{2}$. In this case, an objective A is less likely to pass on his positive signal if the prior beliefs are sufficiently extreme. That is, he may only report $m_A = 1$ with a small probability even if $s_A = 1$, which makes it less likely for objective B to report $m_B = 1$ to the decisionmaker as well even if $s_B = 1$. This increases the reputational cost and decreases the marginal benefit of agenda pushing for biased B, who becomes less likely to report $m_B = 1$. This effect may be very pronounced if there are many agents: sometimes the only equilibrium is an uninformative one in which nothing useful reaches the decisionmaker.

6 Conclusion

Recommendations from well-informed experts often influence the decisionmaker's action. However, even though independent judgment is a sign of objectivity, biased experts selectively incorporate inferior information from potentially biased sources before making their recommendations. Moreover, the intermediary's truth-telling incentives always decrease in those of the source's.

When both the source and the intermediary have very low levels of reputational concerns, any policy measure making it more costly for the source to lie may reduce the quality of the intermediary's recommendation to such an extent that the decisionmaker is strictly worse off. A more effective measure is to target the intermediary instead. Moreover, when the source and the intermediary are sufficiently concerned about their reputations, policy measures encouraging independent reporting by severing the ties between them are more effective than making it more costly for either to lie under indirect communication. The source prefers communicating through an intermediary if both have relatively high reputational concerns, because his message influences the decisionmaker at a negligible reputation cost. But a very biased source may avoid an overly biased intermediary.

Appendix

Proof of Proposition 1: Given A's strategy described in the text, biased B chooses a message m_B to maximize his expected payoff given his private information m_A, s_B. Let $\hat{\eta}_1^B \equiv Pr(\eta = 1|m_B = 1)$, $\hat{\eta}_0^B \equiv Pr(\eta = 1|m_B = 0)$. Then, the difference in B's expected payoff between reporting $m_B = 1$ and reporting $m_B = 0$ is:

$$\begin{aligned} & EU_B(m_B = 1|m_A, s_B) - EU_B(m_B = 0|m_A, s_B) \\ = \;& \hat{\eta}_1^B - \hat{\eta}_0^B - \beta Pr(\eta = 0|m_A, s_B)\Big[Pr(B = o|m_B = 0, \eta = 0) - Pr(B = o|m_B = 1, \eta = 0)\Big] \\ - \;& \beta Pr(\eta = 1|m_A, s_B)\Big[Pr(B = o|m_B = 0, \eta = 1) - Pr(B = o|m_B = 1, \eta = 1)\Big]. \end{aligned}$$

Observe that the difference in B's information is captured in B's estimate of the true state $Pr(\eta = 0|m_A, s_B)$: each estimate induces a different linear combination of his posterior objectivities in the eyes of C. These estimates are respectively:

$$\begin{aligned} \Gamma_1 &\equiv Pr(\eta = 0|m_A = 1, s_B = 0) \\ &= \frac{[1 - p_A + p_A(1-\theta_A)(1-x)]p_B}{[1 - p_A + p_A(1-\theta_A)(1-x)]p_B + [p_A + (1-p_A)(1-\theta_A)(1-x)](1-p_B)}; \\ \Gamma_2 &\equiv Pr(\eta = 0|m_A = 0, s_B = 0) = \frac{p_A p_B}{p_A p_B + (1-p_A)(1-p_B)}; \\ \Gamma_3 &\equiv Pr(\eta = 0|m_A = 0, s_B = 1) = \frac{p_A(1-p_B)}{(1-p_B)p_A + p_A(1-p_B)}; \\ \Gamma_4 &\equiv Pr(\eta = 0|m_A = 1, s_B = 1) \\ &= \frac{[1 - p_A + p_A(1-\theta_A)(1-x)](1-p_B)}{[1 - p_A + p_A(1-\theta_A)(1-x)](1-p_B) + [p_A + (1-p_A)(1-\theta_A)(1-x)]p_B}. \end{aligned}$$

Because $p_B > p_A$, simple calculations can show that $\Gamma_2 > \Gamma_1 > \frac{1}{2} > \Gamma_3 > \Gamma_4$.

To simplify notations, denote C's estimates of B's posterior objectivity respectively as:

$$\begin{aligned} \tau_1 &\equiv Pr(B = o|m_B = 1, \eta = 0) \\ &= \frac{(1-p_B)\theta_B}{1 - p_B + (1-\theta_B)p_B\Big[(1-y_1)[1 - p_A + p_A(1-\theta_A)(1-x)] + (1-y_2)p_A[\theta_A + (1-\theta_A)x]\Big]}; \\ \tau_2 &\equiv Pr(B = o|m_B = 0, \eta = 0) \\ &= \frac{\theta_B}{\theta_B + (1-\theta_B)[y_1(1 - p_A + p_A(1-\theta_A)(1-x)) + y_2 p_A(\theta_A + (1-\theta_A)x))]}; \\ \tau_3 &\equiv Pr(B = o|m_B = 0, \eta = 1) \\ &= \frac{\theta_B}{\theta_B + (1-\theta_B)[y_1(p_A + (1-p_A)(1-\theta_A)(1-x)) + y_2(1-p_A)(\theta_A + (1-\theta_A)x))]}; \end{aligned}$$

$$\tau_4 \equiv Pr(B = o|m_B = 1, \eta = 1)$$
$$= \frac{p_B\theta_B}{p_B + (1-\theta_B)(1-p_B)\Big[(1-y_1)[p_A + (1-p_A)(1-\theta_A)(1-x)] + (1-y_2)(1-p_A)[\theta_A + (1-\theta_A)x]\Big]}.$$

Then, for biased B to report truthfully, the following incentive constraints (IC) must be satisfied:

$$\hat{\eta}_1^B - \hat{\eta}_0^B \leq \Delta_1 \equiv \beta[\Gamma_1(\tau_2 - \tau_1) + (1-\Gamma_1)(\tau_3 - \tau_4)]; \qquad (IC_1^B)$$
$$\hat{\eta}_1^B - \hat{\eta}_0^B \leq \Delta_2 \equiv \beta[\Gamma_2(\tau_2 - \tau_1) + (1-\Gamma_2)(\tau_3 - \tau_4)]; \qquad (IC_2^B)$$
$$\hat{\eta}_1^B - \hat{\eta}_0^B \geq \Delta_3 \equiv \beta[\Gamma_3(\tau_2 - \tau_1) + (1-\Gamma_3)(\tau_3 - \tau_4)]; \qquad (IC_3^B)$$
$$\hat{\eta}_1^B - \hat{\eta}_0^B \geq \Delta_4 \equiv \beta[\Gamma_4(\tau_2 - \tau_1) + (1-\Gamma_4)(\tau_3 - \tau_4)]. \qquad (IC_4^B)$$

The first two incentive constraints (IC_1^B and IC_2^B) concern the case when $s_B = 0$; and the last two (IC_3^B and IC_4^B) concern the case when $s_B = 1$. The left hand side (LHS) of the above ICs is the same, which measures B's net benefit from pushing his agenda. The right hand side (RHS) of the above ICs measures B's net reputation cost if he reports $m_B = 1$ versus $m_B = 0$ given m_A and his signal s_B. B's reporting truthfulness depends on how large his net benefit of agenda-pushing is relative to his net reputation cost. Recall from Lemma 1 that $\Gamma_2 > \Gamma_1 > \Gamma_3 > \Gamma_4$, thus we can rank the RHS of these ICs. For instance, $\Delta_2 - \Delta_1 = (\Gamma_2 - \Gamma_1)[\tau_2 - \tau_3 + \tau_4 - \tau_1]$, and other comparisons are similar.

The term $[\tau_2 - \tau_3 + \tau_4 - \tau_1]$ is positive if B's expected posterior reputation of giving correct predictions $(\tau_2 + \tau_4)$ is larger than that of giving wrong predictions $(\tau_3 + \tau_1)$; it is negative otherwise. Simple calculations can show that it is positive but decreasing in p_A at $p_A \approx \frac{1}{2}$, and may be negative at p_A sufficiently high. Because this term is convex in p_A, there exists a cutoff value $\overline{p}_A$ such that for all $p_A \leq \overline{p}_A$, the difference is always positive. Moreover, it is always positive if β is sufficiently high or close to 0; or if θ_B is sufficiently close to 1. In these cases, B is still better off to report $s_B = 1$ truthfully than to report $m_B = 1$ just to appear independent from A, which is valuable to him in this model.

Next, let the probabilities that B reports his signal $s_B = 1$ truthfully be respectively: $z_1 \equiv Pr(m_B = 1|m_A = 0, s_B = 1)$, $z_2 \equiv Pr(m_B = 1|m_A = 1, s_B = 1)$. For example, if B reports both signals truthfully, then $y_1 = y_2 = z_1 = z_2 = 1$. The following lemma shows that biased B never lies if his signal supports his bias. If his signal does not support his bias, then he lies more if A's message does $(m_A = 1)$ than if it does not $(m_A = 0)$.

Lemma 1 *In any (continuation) equilibrium, $z_1 = 1, z_2 = 1$. Also, $y_1 < y_2$ if $y_2 > 0$.*

Proof: given our assumption that A's signal is sufficiently uninformative, by the argument above, $[\tau_2 - \tau_3 + \tau_4 - \tau_1] > 0$. Then B's net reputation cost can be ranked such that $\Delta_2 > \Delta_1 > \Delta_3 > \Delta_4$. If none of the ICs binds, then it must be $z_1 = z_2 = 1, y_1 = y_2 = 0$. This occurs if B always want to report $m_B = 1$ regardless of his information.

If any of the constraints binds, there are two possibilities. First, suppose that $0 \le z_1 < 1$, or $0 \le z_2 < 1$, then IC_3^B or IC_4^B must be binding, and $y_1 = y_2 = 1$. This possibility corresponds to the case when B reports $s_B = 0$ truthfully but lies with positive probability when $s_B = 1$. As a result, $m_B = 1$ implies $s_B = 1$, and $\hat{\eta}_1^B = p_B$, thus the net benefit for B to send $m_B = 1$ versus $m_B = 0$ is positive ($\hat{\eta}_1^B - \hat{\eta}_0^B > 0$). On the reputation side, however, we can show that $\tau_2 < \tau_1$ and $\tau_3 < \tau_4$. That is, B's net reputation cost is negative. Together, B is strictly better off reporting $m_B = 1$, thus he will deviate and do so, which is a contradiction. Intuitively, in this putative equilibrium, $m_B = 1$ results from signal $s_B = 1$ *and* from an objective B, and as such is both credible and also a good sign of objectivity. The second possibility is if $z_1 = 1, z_2 = 1$, $0 \le y_1 < y_2$. It can be checked that $[\tau_2 - \tau_3 + \tau_4 - \tau_1] > 0$ in this case, thus it is a possible (continuation) equilibrium. □

Given Lemma 1, we know that in equilibrium $z_1 = z_2 = 1$. Also, the only three possible continuation equilibria are: $y_1 = y_2 = 0$, which is the case if neither IC_B^1 nor IC_B^2 holds; $y_2 > 0, y_1 = 0$ (it occurs when IC_B^2 binds) or $y_1 > 0, y_2 = 1$ (it occurs when IC_B^1 binds). We now proceed to characterize the continuation equilibria.

First, consider $\hat{\eta}_1^B - \hat{\eta}_0^B$, the possible gain in term of agenda pushing for a biased B, given his strategy. The inverses of $\hat{\eta}_1^B, \hat{\eta}_0^B$ are respectively:

$$\frac{1}{\hat{\eta}_1^B} = 1 + \frac{1 - p_B + p_B(1-\theta_B)\Big[(1-y_1) - (y_2 - y_1)p_A(\theta_A + (1-\theta_A)x)\Big]}{p_B + (1-p_B)(1-\theta_B)\Big[(1-y_1) - (y_2-y_1)(1-p_A)(\theta_A + (1-\theta_A)x)\Big]};$$

$$\frac{1}{\hat{\eta}_0^B} = 1 + \frac{p_B\Big[\theta_B + (1-\theta_B)[y_1 + (y_2-y_1)p_A(\theta_A + (1-\theta_A)x)]\Big]}{(1-p_B)\Big[\theta_B + (1-\theta_B)[y_1 + (y_2-y_1)(1-p_A)(\theta_A + (1-\theta_A)x)]\Big]}.$$

Simple calculation can show that at $\hat{\eta}_1^B$ increases in y_1, y_2. But $\hat{\eta}_0^B$ increases in y_1 and decreases in y_2, because B reports $m_B = 0$ with y_1 if $m_A = 1, s_B = 0$. That is, the higher is y_1, the more likely that B does not use A's information, which may result from $s_A = 1$. However, since $\frac{\partial^2 \hat{\eta}_1^B}{\partial {y_1}^2} > 0$, $\frac{\partial^2 \hat{\eta}_0^B}{\partial {y_1}^2} < 0$ and $\frac{\partial \hat{\eta}_1^B}{\partial y_1} - \frac{\partial \hat{\eta}_0^B}{\partial y_1} > 0$ at $y = 0$, B's net benefit from pushing his agenda $\hat{\eta}_1^B - \hat{\eta}_0^B$ increases in y_1 and y_2.

At $y_1 = y_2 = 0$, i.e., if biased B lies completely, $\hat{\eta}_1^B - \hat{\eta}_0^B = \frac{2p_B - 1}{1+(1-\theta_B)} > 0$. Thus B's net benefit from agenda pushing is always positive, because of the presence of the objective B; it is also increasing in B's reporting truthfulness y_1, y_2.

Step 1: No truthful revelation equilibrium. To begin with, if B always reports truthfully ($y_1 = y_2 = z_1 = z_2 = 1$), then $[\tau_2 - \tau_3 + \tau_4 - \tau_1] = 0$, because C does not update her belief about B's objectivity at all. This means that B has zero reputation cost, but his benefit of reporting $m_B = 1$ is positive, thus he will deviate, a contradiction.

Step 2: Total distortion equilibrium. When would biased B lie completely if his signal does not support his bias? At $y_1 = y_2 = 0$, B's highest reputation cost occurs if both pieces of his information are against his bias. At $s_B = 0, m_A = 0$, his reputation cost becomes:

$$\Delta_2 = \beta(1-\theta_B)\Big[\frac{\Gamma_2}{1 - p_B\theta_B} + \frac{1-\Gamma_2}{1-(1-p_B)\theta_B}\Big].$$

It can be shown that if $\beta \leq \frac{2p_B - 1}{2-\theta_B}$, or if $\theta_B \approx 1$, the benefit of lying outweighs B's reputation cost, and he lies completely. When $\hat{\eta}_1^B - \hat{\eta}_0^B \leq \Delta_2$ at $y_1 = y_2 = 0$, however, B's expected reputation cost is too high for him to lie completely, which occurs if β is sufficiently high. For example, a total distortion equilibrium does not exist if $\beta \geq 1$.

Step 3: Strong and weak distortion equilibrium. When $\hat{\eta}_1^B - \hat{\eta}_0^B \leq \Delta_2$ at $y_1 = y_2 = 0$, biased B reports truthfully with some probability. There are two possibilities: either a strong distortion equilibrium ($y_1 = 0, y_2 > 0$), or a weak distortion equilibrium ($y_2 = 1, y_1 > 0$) may occur. To learn which equilibrium occurs, we need to compare B's net benefit of lying $\hat{\eta}_1^B - \hat{\eta}_0^B$ with his net reputation cost Δ_2 for any given x.

Specifically, if B's net benefit is higher than his net cost at $y_1 = 0, y_2 = 1$, then $y_1 = 0, y_2 > 0$ is part of the equilibrium. The mixing probability y_2 is determined by a binding IC_B^2. Otherwise, $y_1 > 0, y_2 = 1$ is part of the equilibrium, and the mixing probability y_1 is determined by a binding IC_B^1. Because B's maximum benefit from lying is bounded away from 1, thus for any given x, the LHS of IC_2^B is smaller than the RHS at $y_1 = 0, y_2 = 1$ for β sufficiently large. Thus there exists a cutoff value $\overline{\beta}$ such that if $\beta \leq \overline{\beta}$, the unique continuation equilibrium involves strong distortion and if $\beta \geq \overline{\beta}$, the unique continuation equilibrium involves weak distortion.

Also, we can see how B reacts to A's truth telling. Note that the relative benefit of lying is increasing in y_1, y_2. Moreover, it varies with x in the following way:

$$sign(\frac{\partial \hat{\eta}_1^B}{\partial x}) = sign\Big((y_2 - y_1)[p_B^2 p_A - (1-p_B)^2(1-p_A) + p_B(1-p_B)(1-\theta_A)(2p_A - 1)(1-y_1)]\Big);$$
$$sign(\frac{\partial \hat{\eta}_0^B}{\partial x}) = sign\Big((y_1 - y_2)(2p_A - 1)[\theta_B + (1-\theta_B)y_1]\Big).$$

Thus when $y_2 > y_1$, which is the case in equilibrium, the gain in agenda pushing increases in x as well. On the other hand, simple algebra can show that $\tau_2 - \tau_1$ and $\tau_3 - \tau_4$ decrease in y_1, y_2 and x. Because Γ_2 does not depend on x, Δ_2 decreases in all three mixing probability: the more truthful A and B are, the less C changes her estimates of their objectivity. The remaining case is when IC_B^1 is binding, in which case we have $\frac{\partial \Gamma_1}{\partial x} < 0$ and $sign(\frac{\partial \Delta_1}{\partial x}) = sign(\frac{\partial \Gamma_1}{\partial x}(\tau_2 - \tau_1 - \tau_3 + \tau_4)) < 0$. Thus the reputation cost Δ_1 decreases in the mixing probabilities as well.

Therefore for a given $x > 0$, if x decreases, the gain in agenda pushing for B (the LHS) decreases while his reputation cost (the RHS) increases. Thus B reports more truthfully in either equilibrium: y_1 or y_2 increases. Biased B's truth-telling incentives strictly decrease in those of A's. □

Proof of Proposition 2: Given the continuation equilibria of B, A's (expected) net benefit from reporting $m_A = 1$ instead of reporting $m_A = 0$ given his signal is respectively:

$$\begin{aligned} & E_{m_B}[Pr(\eta = 1|m_B)|m_A = 1, s_A = 0] - E_{m_B}[Pr(\eta = 1|m_B)|m_A = 0, s_A = 0] \\ = \; & (1-\theta_B)(y_2 - y_1)[p_A p_B + (1-p_A)(1-p_B)](\hat{\eta}_1^B - \hat{\eta}_0^B); \\ & E_{m_B}[Pr(\eta = 1|m_B)|m_A = 1, s_A = 1] - E_{m_B}[Pr(\eta = 1|m_B)|m_A = 0, s_A = 1] \\ = \; & (1-\theta_B)(y_2 - y_1)[p_A(1-p_B) + p_B(1-p_A)](\hat{\eta}_1^B - \hat{\eta}_0^B). \end{aligned}$$

Note that A's message only has influence C if B may change his message based on what he hears. If B is objective ($\theta_B = 1$) or if $y_2 = y_1 = 0$, then A's net benefit from agenda pushing becomes 0.

Next, A's expected posterior objectivity if he receives $s_A = 0$ but sends $m_A = 1$ is:

$$\begin{aligned} & E_{\eta, m_B}[Pr(A = o|m_B, \eta)|m_A = 1, s_A = 0] \\ = \; & (\theta_B + (1-\theta_B)y_1)[p_A p_B Pr(A = o|m_B = 0, \eta = 0) + (1-p_A)(1-p_B)Pr(A = o|m_B = 0, \eta = 1)] \\ + \; & (1-\theta_B)(1-y_1)[p_A p_B Pr(A = o|m_B = 1, \eta = 0) + (1-p_A)(1-p_B)Pr(A = o|m_B = 1, \eta = 1)] \\ + \; & p_A(1-p_B)Pr(A = o|m_B = 1, \eta = 0) + p_B(1-p_A)Pr(A = o|m_B = 1, \eta = 1)]. \end{aligned}$$

And his expected posterior objectivity if he reports $m_A = 0$ can be similarly derived. Note that the last part is the effect on A's reputation when B sends $m_B = 1$ because $s_B = 1$.

Similar to before, the net difference in A's expected payoff if he sends $m_A = 1$ versus $m_A = 0$ for a given signal s_A is crucial for his truthful reporting. Let $\kappa = (1-\theta_B)(y_2 - y_1)$, then for biased A to report truthfully, the following two incentive constraints must hold. They respectively require that A is willing to report $s_A = 0$ and $s_A = 1$ truthfully:

$$\begin{aligned}
\kappa(\hat{\eta}_1^B - \hat{\eta}_0^B) \;&\leq\; \frac{\kappa\alpha p_A p_B}{p_A p_B + (1-p_A)(1-p_B)}[Pr(A=o|m_B=0,\eta=0) - Pr(A=o|m_B=1,\eta=0)] \\
&+\; \frac{\kappa\alpha(1-p_A)(1-p_B)}{p_A p_B + (1-p_A)(1-p_B)}[Pr(A=o|m_B=0,\eta=1) - Pr(A=o|m_B=1,\eta=1)]; \quad (1) \\
\kappa(\hat{\eta}_1^B - \hat{\eta}_0^B) \;&\geq\; \frac{\kappa\alpha(1-p_A)p_B}{p_A(1-p_B) + p_B(1-p_A)}[Pr(A=o|m_B=0,\eta=0) - Pr(A=o|m_B=1,\eta=0)] \\
&+\; \frac{\kappa\alpha p_A(1-p_B)}{p_A(1-p_B) + p_B(1-p_A)}[Pr(A=o|m_B=0,\eta=1) - Pr(A=o|m_B=1,\eta=1)]. \quad (2)
\end{aligned}$$

Next, A is also concerned about how objective C thinks he is after each possible message from B and the realized true state. Using Bayes' rule, we have:

$$Pr(A=o|m_B=0,\eta=0) = \frac{\theta_A[\theta_B + p_A(1-\theta_B)y_2 + (1-p_A)(1-\theta_B)y_1]}{p_A(\theta_A+(1-\theta_A)x)(\theta_B+(1-\theta_B)y_2) + [1-p_A(\theta_A+(1-\theta_A)x)](\theta_B+(1-\theta_B)y_1)};$$

$$Pr(A=o|m_B=1,\eta=0) = \frac{\theta_A[1-p_B+p_Ap_B(1-\theta_B)(1-y_2)+(1-p_A)p_B(1-\theta_B)(1-y_1)]}{1-p_B+p_B(1-\theta_B)[p_A(\theta_A+(1-\theta_A)x)(1-y_2)+(1-p_A(\theta_A+(1-\theta_A)x))(1-y_1)]};$$

$$Pr(A=o|m_B=0,\eta=1) = \frac{\theta_A[\theta_B+(1-p_A)(1-\theta_B)y_2+p_A(1-\theta_B)y_1]}{(1-p_A)(\theta_A+(1-\theta_A)x)(\theta_B+(1-\theta_B)y_2)+[1-(1-p_A)(\theta_A+(1-\theta_A)x)](\theta_B+(1-\theta_B)y_1)};$$

$$Pr(A=o|m_B=1,\eta=1) = \frac{\theta_A[p_B+(1-p_A)(1-p_B)(1-\theta_B)(1-y_2)+p_A(1-p_B)(1-\theta_B)(1-y_1)]}{p_B+(1-p_B)(1-\theta_B)[(1-p_A)(\theta_A+(1-\theta_A)x)(1-y_2)+(1-(1-p_A)(\theta_A+(1-\theta_A)x))(1-y_1)]}.$$

Now we can compare the difference in A's expected payoffs if $s_A = 0$ instead of $s_A = 1$. Simple calculations show that first, $EU_A(m_A = 1|s_A = 0) - EU_A(m_A = 0|s_A = 0) > EU_A(m_A = 1|s_A = 1) - EU_A(m_A = 0|s_A = 1)$, or whenever A prefers reporting $m_A = 1$ if $s_A = 0$, he strictly prefers reporting $m_A = 1$ if $s_A = 1$. Intuitively, A has less to lose in term of reputation if $s_A = 1$, because it is more likely that the true state $\eta = 1$ if $s_A = 1$, thus B is more likely to report $m_B = 1$ correctly, which is a less negative sign of A's objectivity. Second, the difference in A's posterior objectivity, $Pr(A = o|m_B = 0, \eta) - Pr(A = o|m_B = 1, \eta)$, decreases in y_1 but increases in y_2. Also, because A's message only affect that of B's when $s_B = 0$, and Proposition 2 shows that $y_2 > y_1$ if $y_2 > 0$, $m_B = 0$ is a better signal of A's objectivity than $m_B = 1$.

To know A's equilibrium behavior if $s_A = 0$, first recall from Proposition 1 that if $\beta \leq \frac{2p_B-1}{1+(1-\theta_B)}$ or if $\theta_B \approx 1$, B always reports $m_B = 1$ $(y_1 = y_2 = 0)$. In this case, A's message has no influence at all on C, thus he can report in any way: $x \in [0, 1]$.

Second, if $y_2 > 0$, A influences B by changing his message with some probability. Simple calculations can show that A's net reputation cost (the RHS of IC (1)) is proportional to $\alpha\theta_A(1 - \theta_A)(1 - x)(1 - \theta_B)(y_2 - y_1)$. Clearly, A reports $x = 0$ if his net reputation cost is approximately 0, which is strictly less than the positive net agenda-pushing effectiveness $\hat{\eta}_1^B - \hat{\eta}_0^B$ on the LHS. This may occur in three cases. First, because the LHS of IC (1) is strictly positive, there exists a cutoff value $\underline{\alpha}$ such that IC (1) binds at $x = 0, y_1 = 0, y_2 = 1$. If $\alpha \leq \underline{\alpha}$, then IC (1) fails to hold for any $x > 0$ because the net benefit for A to lie to push his agenda is higher than the maximum reputation cost he has to pay; if $\alpha \geq \underline{\alpha}$ instead, $x > 0$ for at least some value of β. Next, this is also the case when A's perceived objectivity is not responsive to his messages due to his extreme prior objectivity ($\theta_A \approx 0$ or 1). Finally, it may occur because A has negligible influence on B $(y_1 \approx y_2)$. Note that the last case is true if either $y_2 \approx 0$, which occurs when β is sufficiently low, or if $y_1 \approx 1$, which occurs when β is sufficiently high, A lies completely because his impact on B, and hence his perceived objectivity is negligible.

Third, recall that if β is sufficiently high, one of B's incentive constraints IC_B^1 or IC_B^2 binds (Proposition 2 shows that y_1, y_2 cannot both be strictly between 0 and 1). Moreover, a binding IC_B^1 implicitly defines B's one possible best response to x: $y_1^{BR}(x),\ y_2 = 1$. Similarly, a binding IC_B^2 implicitly define the other possible best response of B to x: $y_2^{BR}(x),\ y_1 = 0$.

We know from the above that a binding IC (1) implicitly define A's best response $BR(y_1, y_2)$. Moreover, $x^{BR}(0, 0) = 0$ and $x^{BR}(0, 1) = 0$ if $\alpha \leq \underline{\alpha}$. Because x decreases in y_1 and y_1 decreases in x, $x^{BR}(y_1, 1) = 0$ for any y_1. Intuitively, if α is sufficiently low, A always reports $m_A = 1$ in equilibrium. If $\alpha \geq \underline{\alpha}$, however, $x^{BR}(0, 1) > 0$. Then if $y_2^{BR}(1) < 1$, because y_2 decreases in x, $y_2^{BR}(0) > y_2^{BR}(1) \geq 0$. As mixing probabilities, $x, y_2 \in [0, 1]$. Thus by the intermediate value theorem, the two best responses intersect, and there exists a strong distortion equilibrium in which $x \in [0, 1], y_2 \in (0, 1)$.

If, however, $y_2^{BR}(1) \geq 1$, then no strong distortion equilibrium exists, which is the case if IC_B^2 does not hold at $x = 1, y_1 = 0, y_2 = 1$. Moreover, $y_2^{BR}(1) \geq 1$ implies that $y_1^{BR}(1) > 0$. Because y_1 decreases in x, $0 < y_1^{BR}(1) < y_1^{BR}(0) < 1$. We know that $x^{BR}(0,1) > 0$ and $x^{BR}(y_1, 1) = 0$ because if $y_1 \approx 1$, A's reputation cost (the RHS of IC (1)) is approximately 0, strictly smaller than the RHS. Because x decreases in y_1, and both $x, y_1 \in [0, 1)$, these two best responses intersect and there exists a weak distortion equilibrium in which $x \in [0, 1), y_1 \in (0, 1)$.

Finally, if α is sufficiently high and if $y_2^{BR}(0) > 1$ and $y_2^{BR}(1) < 1$, there may exist multiple equilibria in which either $x_1 > 0, y_2 > 0$ or $x_2 > 0, y_1 > 0$. To see this, note from biased B's incentive constraint IC_B^2 in the proof of Proposition 1 that if $y_2^{BR}(0) \leq 1$, then for any x, only a strong distortion equilibrium may exist; if $y_2^{BR}(1) \geq 1$ instead, then for any value of x, only a weak distortion equilibrium exists. In particular, let β_1 be the value such that IC_B^2 binds at $x = 0, y_1 = 0, y_2 = 1$ and $\overline{\beta}$ be the value such that IC_B^2 binds at $x = 1, y_1 = 0, y_2 = 1$. Then multiple equilibria may occur only if $\beta \in [\beta_1, \overline{\beta})$. Intuitively, in this range, B's behavior is very sensitive to A's behavior. Thus if $m_A = 1$ is very credible, for instance if B's signal quality is not too much higher than that of A's, or if A has very high reputational concerns, B can rely more on A's message and lie more in a strong distortion equilibrium. But if $m_A = 1$ is not very credible, then B prefers relying on his own signal more, which gives rise to a weak distortion equilibrium. □

Proof of Proposition 3: From the proof of Proposition 1 and Proposition 2, we know that for $p_A \leq \overline{p}_A$, in equilibrium the agents lie against their signals $s_A = 0$ and $s_B = 0$. In particular, if $\alpha \geq \underline{\alpha}$, $\beta \geq \underline{\beta}$ but not too large (β sufficiently close to $\overline{\beta}$, there exists an interior equilibrium in which the mixing probabilities $x, y_1 \in (0, 1)$ (or $x, y_2 \in (0, 1)$). Moreover, y_1 decreases in x; and x decreases in y_1. Thus x, y_1 are substitutes: whenever one agent reports more truthfully, the other becomes less truthful.

As shown in Proposition 1, y_2 increases in x. However, A's best response with respect to y_2 is ambiguous because both his net agenda pushing effectiveness and his net reputation cost increase in y_2. However, if p_B is sufficiently high, A's equilibrium mixing condition (1) can be used to show that A's agenda-pushing effectiveness (the LHS of (1)) increases less in y_2 than his reputation cost (the RHS of (1)). Thus A reports more truthfully as B does: x increases in y_2 if p_B is sufficiently high. Intuitively, because biased B follows A's message in a strong distortion equilibrium, a wrong message is more likely due to A's distortion if B's signal is very precise. □

Proof of Proposition 4: Let the weights on the agents' reputation be α_i, α_j and their signal qualities be p_i, p_j respectively. Suppose they report signals $s_i = 0$ and $s_j = 0$ with probabilities be x_i, x_j respectively. Recall from the text that here biased $i = A, B$ maximizes:

$$E_{m_j}[Pr(\eta = 1|m_i, m_j)|s_i] + E_\eta[Pr(i = o|m_i, \eta)|s_i].$$

To simplify notations, denote C's belief that the true state is 1, $Pr(\eta = 1|m_i = 0, m_j = 0)$, as $Pr(1|0,0)$; other beliefs are similarly denoted. Observe that agent j's message does not affect C's posterior estimate of agent i's objectivity because she evaluates i after observing the true state. However, m_j affects the marginal impact of i's message on C's action.

In particular, for biased i to report truthfully, the following two incentive constraints must hold:

$$\begin{aligned} & Pr(m_j = 1|s_i = 0)[Pr(1|1,1) - Pr(1|0,1)] + Pr(m_j = 0|s_i = 0)[Pr(1|1,0) - Pr(1|0,0)] \\ \geq\ & \alpha_i \sum_\eta Pr(\eta|s_i = 0)[Pr(i = o|m_i = 1, \eta) - Pr(i = o|m_i = 0, \eta)]; \qquad (3) \\ & Pr(m_j = 1|s_i = 1)[Pr(1|1,1) - Pr(1|0,1)] + Pr(m_j = 0|s_i = 1)[Pr(1|1,0) - Pr(1|0,0)] \\ \geq\ & \alpha_i \sum_\eta Pr(\eta|s_i = 1)[Pr(i = o|m_i = 1, \eta) - Pr(i = o|m_i = 0, \eta)]. \qquad (4) \end{aligned}$$

Given the agents' strategies, C's action after hearing both messages is:

$$\begin{aligned} Pr(1|0,0) &= \frac{(1-p_i)(1-p_j)}{p_i p_j + (1-p_i)(1-p_j)}; \\ Pr(1|0,1) &= \frac{(1-p_i)[1-(1-p_j)(\theta_j + (1-\theta_j)x_j)]}{(1-p_i)[1-(1-p_j)(\theta_j + (1-\theta_j)x_j)] + p_i[1-p_j(\theta_j + (1-\theta_j)x_j)]}; \\ Pr(1|1,0) &= \frac{(1-p_j)[1-(1-p_i)(\theta_i + (1-\theta_i)x_i)]}{(1-p_j)[1-(1-p_i)(\theta_i + (1-\theta_i)x_i)] + p_j[1-p_i(\theta_i + (1-\theta_i)x_i)]}; \\ Pr(1|1,1) &= \frac{1}{1 + \frac{[1-p_j(\theta_j + (1-\theta_j)x_j)]}{[1-(1-p_j)(\theta_j + (1-\theta_j)x_j)]} \cdot \frac{[1-p_i(\theta_i + (1-\theta_i)x_i)]}{[1-(1-p_i)(\theta_i + (1-\theta_i)x_i)]}}. \end{aligned}$$

Moreover, because of the presence of objective agent, it is simple to show that $Pr(1|1,1) > Pr(1|0,1)$ and $Pr(1|1,0) > Pr(1|0,0)$. Next, it can also be shown that the difference $Pr(1|1,1) - Pr(1|0,1) - [Pr(1|1,0) - Pr(1|0,0)] > 0$. Intuitively, $m_i = 1$ has a higher marginal impact in term of agenda pushing if m_j supports rather than contradicts i's message. The reason is that relative to the case with only objective agents, biased i is more likely to distort when $\eta = 0$ than when $\eta = 1$, in which case his signal is more likely to be $s_i = 1$ and he does not distort. Therefore if $m_j = 0$, the decisionmaker infers that it is more likely that $\eta = 0$ and $m_i = 1$ is more likely to be a result of distortion than if $m_j = 1$.

This also implies that if IC (3) binds or fails to hold, IC (4) holds strictly. The reason is that if $s_i = 1$, it is more likely that $s_j = 1$ as well, thus the agenda pushing effectiveness is higher for i to report $m_j = 1$ as opposed to $m_i = 0$. On the other hand, the net reputation cost is smaller if $s_i = 1$ because his message $m_i = 1$ is more likely to be correct. Thus i always reports $s_i = 1$ truthfully.

Next, differentiate i's agenda pushing effectiveness (the LHS of IC (3)) with respect to x_j, and we have, up to a factor $(1-\theta_j)$:

$$\begin{aligned} & Pr(m_j = 1|s_i = 0)[\frac{\partial}{x_j}(Pr(1|1,1)) - \frac{\partial}{x_j}(Pr(1|0,1))] \\ - \quad & [p_i p_j + (1-p_i)(1-p_j)][Pr(1|1,1) - Pr(1|0,1) - [Pr(1|1,0) - Pr(1|0,0)]. \end{aligned}$$

From the discussion above, the second line is clearly negative. Moreover, i's net benefit from lying is decreasing in x_j because $\frac{\partial}{\partial x_j}Pr(1|1,1) < \frac{\partial}{\partial x_j}Pr(1|0,1)$. Therefore x_i increases in x_j. The biased agents' truth-telling probabilities are complements because if one becomes more truthful, the other has a smaller (expected) marginal impact on C, thus he lies less as well. □

Proof of Proposition 5: Consider independent reporting first. Suppose that in equilibrium, A reports $s_A = 0$ with positive probability x_A, then his ex ante expected payoff becomes:

$$\begin{aligned} EU_A^1 &\equiv E_{s_A}\Big[E_{m_B}[Pr(\eta = 1|m_A, m_B)|s_A] + \alpha E_\eta[Pr(A = o|m_A, \eta)|s_A]\Big] \\ &= \frac{1}{2}[(1 + (1-\theta_B)(1-x_B))Pr(1|1,1) + (\theta_B + (1-\theta_B)x_B)Pr(1|1,0)] \\ &+ \alpha Pr(A = o|m_A = 1). \end{aligned} \tag{5}$$

Intuitively, because A is indifferent between reporting $m_A = 1$ or $m_A = 0$ if $s_A = 0$, while he strictly prefers reporting $m_A = 1$ if $s_A = 1$, his expected payoff is the same as if he always reports $m_A = 1$. Moreover, biased A's ex ante payoff can be compared with a sum of C's prior beliefs $\frac{1}{2} + \alpha\theta_A$:

$$\begin{aligned} & EU_A^1 - (\frac{1}{2} + \alpha\theta_A) \qquad (6) \\ = \quad & EU_A^1 - \sum_{l,k\in\{0,1\}} Pr(m_A = l, m_B = k)Pr(\eta = 1|m_A = l, m_B = k) - \sum_{l\in\{0,1\}} \alpha Pr(m_A = l)Pr(A = o|m_A = l) \\ = \quad & Pr(m_A = 0)\Big[Pr(m_B = 1|s_A = 0)[Pr(1|1,1) - Pr(1|0,1)] + Pr(m_B = 0|s_A = 0)[Pr(1|1,0) - Pr(1|0,0)]\Big] \\ - \quad & \alpha Pr(m_A = 0)[Pr(A = o|m_A = 0) - Pr(A = o|m_A = 1)] \\ = \quad & 0. \end{aligned}$$

The first equality is due to the law of iterated expectations, while the last one is due to biased A's equilibrium mixing condition IC (3), as given in Proposition 4. Therefore if biased A cares sufficiently about his reputation to report $m_A = 0$ sometimes, his net gain in term of agenda-pushing effectiveness must exactly be equal to his loss in posterior objectivity.

Second, if A can exert influence through intermediary B ($\max\{y_1, y_2\} > 0$), and suppose that in equilibrium $x > 0$, biased A's ex ante expected payoff from indirect communication is:

$$\begin{aligned} & EU_A^2 \\ = \; & \frac{1}{2}[1 + (1-\theta_B)(1-y_1)][Pr(\eta = 1|m_B = 1) + \alpha Pr(A = o|m_B = 1, \eta = 1) + \alpha Pr(A = o|m_B = 1, \eta = 0)] \\ + \; & \frac{1}{2}[\theta_B + (1-\theta_B)y_1][Pr(\eta = 1|m_B = 0) + \alpha Pr(A = o|m_B = 0, \eta = 1) + \alpha Pr(A = o|m_B = 0, \eta = 0)]. \quad (7) \end{aligned}$$

Similar argument can show that $EU_A^2 = \frac{1}{2} + \alpha\theta_A$ if there exists a fully mixed equilibrium where both A and B report signals against their agenda sometimes. Thus if $x_A > 0, x > 0, \max\{y_1, y_2\} > 0$, A is indifferent in ex ante terms between these two ways of communication.

Next, if biased A strictly prefers lying, then his ex ante expected payoff is strictly higher than the prior: $\frac{1}{2} + \alpha\theta_A$. For instance, suppose that $x_A = 0$. Then A's reputational cost is so low that $EU_A^1(m_A = 1|s_A = 0) > EU_A^1(m_A = 0|s_A = 0)$ at $x_A = 0$, thus he is worse off if he reports truthfully with any infinitesimally small probability. To see this, note that Expression (6) is strictly positive because A's mixing constraint IC (3) no long binds.

Recall from Proposition 4 that A always reports $m_A = 1$ if α is sufficiently low or if θ_A is sufficiently high. A simple comparison of A's ex ante expected payoffs can show that if β is sufficiently low that $y_2 \approx 0$, then A prefers independent reporting if θ_B is sufficiently close to 0. Formally, this is because $Pr(\eta = 1|1, 1) > Pr(\eta = 1|m_B = 1)$. Otherwise, A prefers communicating through an intermediary.

Finally, if β is sufficiently high, biased A always reports $m_A = 1$ with indirect communication, regardless of his own reputational concerns. Therefore A with moderately high reputational concerns — who would report $s_A = 0$ truthfully with some probability if he sends his own message — always prefers indirect communication if the intermediary is very concerned about his reputation. □

References

Austen-Smith, D. (1990): "Information Transmission in Debate," *American Journal of Political Science*, 34(1), 124–152.

Avorn, J., M. Chen, and R. Hartley (1982): "Scientific Versus Commercial Sources of Influence on the Prescribing Behavior of Physicians," *American Journal of Medicine*, 73(1), 4–8.

Bandler, J. (2005): "Advice for Sale: How Companies Pay TV Experts For On-Air Product Mentions," The Wall Street Journal, April 19, 2005.

Brennan, T., D. Rothman, L. Blank, D. Blumenthal, S. Chimonas, J. Cohen, J. Golden, J. Kassirer, H. Kimball, J. Naughton, and N. H. Smelser (2006): "Health industry practices that create conflicts of interest: a policy proposal for academic medical centers," *Journal of the American Medical Association*, 295(4), 429–433.

Carey, B. (2006): "Study Finds a Link of Drug Makers to Psychiatrists," The New York Times.

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

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.

Durbin, E., and G. Iyer (2006): "Corruptible Advice," Mimeo, University of California, Berkeley.

Inderst, R., and M. Ottaviani (2007): "(Mis)selling through Agents," Mimeo, London Business School.

Li, W. (2007a): "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.

——— (2007b): "Peddling Influence through Intermediaries: Propaganda," mimeo, UC Riverside.

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

NYT (2002): "Whistle-blower Says Marketers Broke the Rules to Push a Drug," The New York Times.

——— (2006): "Wal-Mart Enlists Bloggers in P.R. Campaign," The New York Times.

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

Saul, S. (2006): "Unease on Industry's Role in Hypertension Debate," The New York Times.

Watkins, C., L. Moore, I. Harvey, P. Carthy, E. Robinson, and R. Brawn (2003): "Characteristics of General Practitioners Who Frequently See Drug Industry Representatives: One National Cross Sectional Survey," *British Medical Journal*, 326(7400), 1178–9.