Opinions as Incentives*

Yeon-Koo Che[†] Navin Kartik[‡]

February 6, 2008

Abstract

We study a model where a decision maker (DM) must select an adviser to advise her about an unknown state of the world. There is a pool of available advisers who all have the same underlying preferences as the DM; they differ, however, in their prior beliefs about the state, which we interpret as differences of opinion. We derive a tradeoff faced by the DM: an adviser with a greater difference of opinion has greater incentives to acquire information, but reveals less of any information she acquires, via strategic disclosure. Nevertheless, it is optimal to choose an adviser with at least some difference of opinion. The analysis reveals two novel incentives for an agent to acquire information: a "persuasion" motive and a motive to "avoid prejudice." Delegation is costly for the DM because it eliminates both of these incentives. We also study the relationship between difference of opinion and difference of preference.

*We would like to acknowledge the input of Jimmy Chan at early stages of this project. We thank Nageeb Ali, Oliver Board, Arnaud Costinot, Vince Crawford, Jean Guillaume Forend, Michihiro Kandori, Kohei Kawamura, Li Hao, Bart Lipman, Eric Maskin, Carolyn Pitchik, Jennifer Reinganum, Mike Riordan, Ed Schlee, Joel Sobel, and various seminar audiences for their opinions. David Eil, Chulyoung Kim, and Uliana Loginova provided excellent research assistance. Kartik is grateful to the Institute for Advanced Study at Princeton and the National Science Foundation for funding; he also thanks the Institute for its hospitality.

[†]Economics Department, Columbia University; yc2271@columbia.edu.

[‡]Department of Economics, University of California, San Diego; nkartik@ucsd.edu.

"Difference of opinion leads to enquiry." — Thomas Jefferson

1 Introduction

To an average 17th century (geocentric) person, the emerging idea of the earth moving defied common sense. If the earth revolves, then "why would heavy bodies falling down from on high go by a straight and vertical line to the surface of the earth... [and] not travel, being carried by the whirling earth, many hundreds of yards to the east?" (Galilei, 1953, p. 126) In the face of this seemingly irrefutable argument, Galileo Galilei told a famous story, via his protagonist *Salviati* in *Dialogue Concerning the Two Chief World Systems*, about how an observer locked inside a boat, sailing at a constant speed without rocking, cannot tell whether the boat is moving or not. This story, meant to persuade critics of heliocentrism, became a visionary insight now known as the *Galilean Principle of Relativity*.

The above example dramatically illustrates how a different view of the world (literally) might lead to an extraordinary discovery. But the theme it captures is hardly unique. Indeed, difference of opinion is valued in many organizations and situations. Corporations seek diversity in their workforce allegedly to tap creative ideas. Academic research thrives on the pitting of opposing hypotheses. Government policy failures are sometimes blamed on the lack of a dissenting voice in the cabinet, a phenomenon coined "groupthink" by psychologists (e.g. Janis, 1972). Debates between individuals can be more illuminating when they take different views; in their absence, debaters often create an artificial difference by playing "devil's advocate."

Difference of opinion would be obviously valuable if it inherently entails a productive advantage in the sense of bringing new ideas or insights that would otherwise be unavailable. But could it be valuable even when it brings no direct productive advantage? Moreover, are there any costs of people having differing opinions? This paper explores these questions by examining incentive implications of difference of opinion.

We develop a model in which a decision maker, or DM for short, consults an adviser before making a decision. There is an unknown state of the world that affects both individuals' payoff from the decision. We model the DM's decision and the state of the world as real numbers, and assume the DM's optimal decision coincides with the state. Initially, neither the DM nor the adviser has any information about the state beyond their prior views. The adviser can exert effort to try and produce an informative signal about the state, which occurs with probability that is increasing in his effort. The signal could take the form of scientific evidence obtainable by conducting an experiment, witnesses or documents locatable by investigation, a mathematical proof, or a convincing insight that can reveal something about the state. Effort is unverifiable, however, and higher effort imposes a greater cost on the adviser. After the adviser privately observes the information, he strategically communicates with the DM. Communication takes the form of verifiable disclosure: sending a message is costless, but the adviser cannot falsify information, or equivalently, the DM can judge objectively what a signal means. The adviser can, nevertheless, choose not to disclose the information he acquires. Finally, the DM takes her decision optimally given her updated beliefs after communication with the adviser.

This framework captures common situations encountered by many organizations. For instance, managers solicit information from employees; political leaders seek the opinion of their cabinet members; scientific boards consult experts; and journal editors rely on referees. But the model permits broader interpretations: the DM could be the general public (such as 17th century intelligent laymen), and its decision is simply the posterior belief on some matter. In turn, the adviser could be a scientist (such as Galileo), investigator, special counsel, a lobbying group, or a debater trying to sway that belief.

It is often the case, as in the examples mentioned above, that an adviser is interested in the decision taken by DM. We assume initially that the adviser has the same fundamental preferences as the DM about which decision to take in each state, but that he may have a difference of opinion about what the unknown state is likely to be. More precisely, the adviser may disagree with the DM about the prior probability distribution of the unknown state, and this disagreement is common knowledge.[1] Such disagreements abound in many circumstances, as has also been argued by, for example, Banerjee and Somanathan (2001). Consider a firm that must decide which of two technologies to invest in. All employees share the common goal of investing in the better technology, but no one knows which this is. Different employees may hold different beliefs about the viability of each technology, leading to open disagreements about where to invest. As another example, a general and her advisers may agree on the objective of winning a war at minimum cost. They may have

[1] That is, they "agree to disagree." Such an open disagreement may arise from various sources: individuals may simply be endowed with different prior beliefs (just as they may be endowed with different preferences), or they may update certain kinds of public information differently based on psychological, cultural, or other factors (Tversky and Kahneman, 1974; Aumann, 1976; Acemoglu, Chernozhukov, and Yildiz, 2007). Whatever the reason, open disagreements do exist and often persist even after extensive debates and communication.

different beliefs, however, about the strength of the opposition troops, leading to disagreements about how many of their own troops should be sent into combat—disagreements that do not change even when told each other's views. Many political disagreements also seem best viewed through the lens of different prior beliefs rather than different fundamental preferences (Dixit and Weibull, 2007).[2]

Specifically, we model the adviser's opinion as the mean of his (subjective) prior about the state, normalizing the DM's opinion to mean zero. We suppose that there is a rich pool of possible advisers in terms of their opinion, and advisers are differentiated only by their opinion, meaning that a difference of opinion does not come with better ability or lower cost of acquiring information. This formulation allows us to examine directly whether difference of opinion alone can be valuable to the DM, even without any direct productive benefits.[3]

Our main results concern a tradeoff associated with difference of opinion. To see this, suppose first that effort is not a choice variable for the adviser. In this case, the DM has no reason to prefer an adviser with a differing opinion. In fact, unless the signal is perfectly informative about the state, the DM will strictly prefer a like-minded adviser—i.e., one with the same opinion as she has. This is because agents with different opinions, despite having the same preference, will generally arrive at different posteriors about what the right decision is given partially-informative signals. Consequently, an adviser with a differing opinion will typically withhold some information from the DM. This strategic withholding of information entails a welfare loss for the DM, whereas no such loss will arise if the adviser is like-minded.

When effort is endogenous, the DM is also concerned with the adviser's incentive to exert effort; all else equal, she would prefer an adviser who will exert as much effort as possible. We find that differences of opinion provide incentives for information acquisition, for two distinct reasons. First, *an adviser with a difference of opinion is motivated to persuade the DM*. Such an adviser believes that the DM's opinion is wrong, and that by acquiring a signal, he is likely to move the DM's decision towards what he perceives to be the right decision. This motive does not exist for the like-minded adviser. Second, and more subtle, *an adviser with difference of opinion will exert effort to avoid "prejudice."* Intuitively,

[2]To mention just two examples, consider disagreements about how serious the global warming problem is (if it exists all, to some) and how to protect a country against terrorism.

[3]As previously noted, individuals with different backgrounds and experiences are also likely to bring different approaches and solutions to a problem, which may directly improve the technology of production. We abstract from these in order to focus on the incentive implications of difference of opinion.

in equilibrium, an adviser withholds information that is contrary to his opinion, for such information will cause the DM to take an action that the adviser dislikes. Recognizing this, the DM discounts the advice she receives and chooses an action contrary to the adviser's opinion, unless the advice is corroborated by a hard evidence—this equilibrium feature of strategic interaction is what we call a "prejudicial effect." Hence, an adviser with difference of opinion has incentives to seek out information in order to avoid prejudice, a motive that does not exist for a like-minded adviser.

In summary, we find that *difference of opinion entails a loss of information through strategic communication, but creates incentives for information acquisition*. This tradeoff resonates with common notions that, on the one hand, diversity of opinion causes increased conflict because it becomes harder to agree on solutions—this emerges in our analysis as worsened communication; on the other hand (as was recognized by Jefferson, quoted in our epigraph) it also leads to increased efforts to understand and convince other individuals—this emerges here as increased information acquisition.

How should the DM resolve this tradeoff between information acquisition and transmission? We find that *the DM prefers an adviser with some difference of opinion to a perfectly like-minded one*. The reason is that an adviser with sufficiently small difference of opinion engages in only a negligible amount of strategic withholding of information, so the loss associated with such a difference is negligible. By the same token, the prejudicial effect and its beneficial impact on information acquisition is also negligible when the difference of opinion is small. In contrast, the persuasion motive that even a slight difference of opinion generates—and thus the benefit the DM enjoys from its impact on increased effort—is non-negligible by comparison. Therefore, the DM strictly benefits from an adviser with at least a little difference in opinion, and would not optimally choose a like-minded adviser from a rich pool of available individuals.

Sections 2–4 formalize our model and the above logic. Section 5 then augments the model to allow the adviser to differ from the DM in both his opinion and his fundamental preferences over decisions. Heterogeneous preferences have a similar effect as difference of opinion on strategic disclosure. But this similarity does *not* extend to the adviser's choice of effort, because the two attributes are fundamentally distinct in terms of how they motivate the adviser. While an adviser with difference of opinion has a persuasion motive for acquiring a signal—he expects to systematically shift the DM's decision closer to his preferred decision—an adviser with only a difference of preference has no such expectation, and thus has no persuasion motive. For this reason, having an adviser who differs only in preferences yields no clear benefit for the DM.

Nevertheless, we find the difference of preferences to be valuable in the presence of difference of opinion. In other words, an adviser with a different opinion has more incentive to acquire information if he also has an preference bias in the direction congruent to his opinion. This complementarity between preference and opinion implies that the incentive effect on information acquisition will be larger when the adviser is a *zealot*—one who believes that evidence is likely to move the DM's action in the direction of his preference bias—than when he is a *skeptic*—one who is doubtful that information about the state of the world will support his preference bias.

We explore some other issues in Section 6. Of particular interest, we find that the benefit from difference of opinion is lost when the DM delegates the decision authority to the adviser. This observation sheds new light on the merit of delegation in organizational settings (cf. Aghion and Tirole, 1997). We also discuss implications of the adviser's perception of the precision of his own information, or his *confidence*, finding that more confident advisers exert more effort.

Our paper builds on the literature on strategic communication, combining elements from the structure of conflicts of interest in Crawford and Sobel (1982) with the verifiable disclosure game first introduced by Grossman (1981) and Milgrom (1981). The key innovation in this regard is that we endogenize information acquisition and focus on the effects of difference of prior beliefs. We postpone a detailed discussion of the related literature to Section 7, after a full development of our model and analysis. Section 8 then concludes by relating some of our insights to applications, and a brief discussion of possible extensions. The Appendix contains omitted proofs.

2 Model

A decision maker (DM) must take a decision, $a \in \mathbb{R}$. The appropriate decision depends on an unknown state of the world, $\omega \in \mathbb{R}$. The DM lacks the necessary expertise, or finds it prohibitively costly, to directly acquire information about the state, but can choose a single adviser from a pool of available agents to advise her.

Prior Beliefs. We allow individuals—potential advisers and the DM—to have different prior beliefs about the state. Specifically, while all individuals know the state is distributed according to a Normal distribution with variance $\sigma_0^2 > 0$, individual i believes the mean of the distribution is μ_i. The prior beliefs of each person are common knowledge.[4] We will refer to an adviser's prior belief as his *opinion* or *type*, even though it is not private information. Two individuals, i and j, have differences of opinion if $\mu_i \neq \mu_j$. Without loss of generality, we normalize the DM's prior to $\mu = 0$. An adviser with $\mu = 0$ is said to be *like-minded*.

Full-information preferences. All players have the same von Neumann-Morgenstern state-dependent payoff from the DM's decision:

$$u_i(a, \omega) := -(a - \omega)^2.$$

Thus, were the state ω known, players would agree on the optimal decision $a = \omega$. In this sense, there is no fundamental preference conflict. We allow for such conflicts in Section 5. The quadratic loss function we use is a common specification in the literature: it captures the substantive notion that decisions are progressively worse the further they are from the true state, and technically, makes the analysis tractable.

Information Acquisition. Regardless of the chosen adviser's type, his investigation technology is the same, described as follows. He chooses the probability that his investigation is successful, $p \in [0, \bar{p}]$, where $\bar{p} < 1$, at a cost $c(p)$. The function $c(\cdot)$ is smooth, $c''(\cdot) > 0$, and satisfies the Inada conditions $c'(0) = 0$ and $c'(p) \to \infty$ as $p \to \bar{p}$. We will interchangeably refer to p as an effort level or a probability.[5] With probability p, the adviser obtains a signal about the state, $s \sim N(\omega, \sigma_1^2)$. That is, the signal is drawn from a *Normal* distribution with mean equal to the true state and variance $\sigma_1^2 > 0$. With complementary probability $1 - p$, he receives no signal, denoted by $\emptyset$. Thus, effort is success-enhancing in the sense of Green and Stokey (2007) and increases information in the sense of Blackwell (1951).

[4] Although game-theoretic models often assume a common prior, referred to as the *Harsanyi Doctrine*, there is a significant and growing literature that analyzes games with heterogenous priors. Spector (2000) and Banerjee and Somanathan (2001) do so in communication models with exogenous information; in other contexts, examples are Harrington (1993), Yildiz (2003), Van den Steen (2005), and Eliaz and Spiegler (2006). For a general discussion about non-common priors, see Morris (1995).

[5] This is justified because our formulation is equivalent to assuming the adviser chooses some effort e at cost $c(e)$, which maps into a probability $p(e)$.

Communication. After privately observing the outcome of his investigation, the chosen adviser strategically discloses information to the DM. The signal s is "hard" or non-falsifiable. Hence, the adviser can only withhold the signal if he has obtained one; if he did not receive a signal, he has no choice to make. The signal may be non-manipulable because there are large penalties against fraud, information is easily verifiable by the DM once received (even though impossible to acquire herself), or information is technologically hard to manipulate.[6]

Timing. The sequence of events is as follows. First, the DM selects an adviser from an available set of adviser types, $[\underline{\mu}, \overline{\mu}]$, where $\underline{\mu} < 0 < \overline{\mu}$. The selected adviser then chooses effort and observes the outcome of his investigation, both unobservable to the DM. In the third stage, the adviser either discloses or withholds any information acquired. Finally, the DM takes a decision.

As this is multi-stage Bayesian game, it is appropriate to solve it using the concept of perfect Bayesian equilibrium (Fudenberg and Tirole, 1991), or for short, *equilibrium* hereafter. We restrict attention to pure strategy equilibria.

2.1 Interim Bias

As a prelude to our analysis, it is useful to identify the players' preferences over decisions when the state is not known. Throughout, we use subscripts DM and A for the decision maker and adviser, respectively. Under the Normality assumptions in our information structure, the signal and state joint distribution can be written, from the perspective of player $i = DM, A$, as

$$\begin{pmatrix} \omega \\ s \end{pmatrix} \sim N\left(\begin{matrix} \mu_i \\ \mu_i \end{matrix}, \begin{pmatrix} \sigma_0^2 & \sigma_0^2 \\ \sigma_0^2 & \sigma_0^2 + \sigma_1^2 \end{pmatrix} \right).$$

Without a signal about the state, the expected utility of player i is maximized by action μ_i. Suppose a signal s is observed. The posterior of player i is that $\omega | s \sim N\left(\rho s + (1-\rho)\mu_i, \tilde{\sigma}^2\right)$,

[6]Our formulation follows, for example, Shin (1998). Alternatively, we could assume that the adviser must make an assertion that the signal lies in some compact set, $\mathcal{S}$, or $\mathbb{R}$, with the only constraint that $s \in \mathcal{S}$, as formulated by Milgrom (1981). In this case, when the signal is not observed, the adviser has to report $\mathbb{R}$. By endowing the DM with a "skeptical posture" (Milgrom and Roberts, 1986) when the adviser claims any set $\mathcal{S} \neq \mathbb{R}$, our analysis can be extended to this setting.

where $\rho := \frac{\sigma_0^2}{\sigma_0^2+\sigma_1^2}$ and $\tilde{\sigma}^2 := \frac{\sigma_0^2\sigma_1^2}{\sigma_0^2+\sigma_1^2}$ (Degroot, 1970).[7] Player $i = DM, A$ therefore has the following expected utility from action a given s:

$$\begin{aligned}\mathbb{E}[u_i(a,\omega)|s,\mu_i] &= -\mathbb{E}[(a-\omega)^2|s,\mu_i] = -(a-\mathbb{E}[\omega|s,\mu_i])^2 - \mathrm{Var}(\omega|s) \\ &= -\left(a - \{\rho s + (1-\rho)\mu_i\}\right)^2 - \tilde{\sigma}^2. \end{aligned} \quad (1)$$

Clearly, the expected utility is maximized by an action $\alpha(s|\mu_i) := \rho s + (1-\rho)\mu_i$, where $\alpha(s|\mu)$ is simply the posterior mean for a player with type μ.

Equation (1) shows that so long as signals are not perfectly informative of the state ($\rho < 1$), differences of opinion generate conflicts in preferred decisions given any signal, even though fundamental preferences agree. Accordingly, we define the *interim bias* as $B(\mu) := (1-\rho)\mu$. This completely captures the difference in the two players' preferences over actions given any signal because $\alpha(s|\mu) = \alpha(s|0) + B(\mu)$. Observe that for any $\mu \neq 0$, $sign(B(\mu)) = sign(\mu)$ but $|B(\mu)| < |\mu|$. Hence, while interim bias persists in the same direction as prior bias, it is of strictly smaller magnitude because information about the state mitigates prior disagreement about the optimal decision. This simple observation turns out to have significant consequences. The magnitude of interim bias depends upon how precise the signal is relative to the prior; differences of opinion matter very little once a signal is acquired if the signal is sufficiently precise, i.e. for any μ, $B(\mu) \to 0$ as $\rho \to 1$ (equivalently, as $\sigma_1^2 \to 0$ or $\sigma_0^2 \to \infty$).

3 Equilibrium Disclosure Behavior

In this section, we analyze the behavior of adviser and DM in the disclosure sub-game.[8] For this purpose, it will be sufficient to focus on the interim bias of the adviser, $B(\mu)$, and the DM's belief about the probability p that the adviser observes a signal.[9] Hence, we take the pair (B, p) as a primitive parameter in this section. Our objective is to characterize the set $S \in \mathbb{R}$ of signals that the adviser withholds and the action $a_\emptyset$ the DM chooses when there

[7]Since $\sigma_0^2 > 0$ and $\sigma_1^2 > 0$, $\rho \in (0,1)$. However, it will be convenient at points to discuss the case of $\rho = 1$; this should be thought of as the limiting case where $\sigma_1^2 = 0$, so that signals are perfectly informative about the state. Similarly for $\rho = 0$.

[8]Strictly speaking, we are abusing terminology in referring to this as a "sub-game," because the DM does not observe the adviser's effort choice, p.

[9]The subsequent analysis will show why the DM's belief about the adviser's effort, rather than the true effort, is what matters for disclosure behavior. (Of course, we will require this belief to be correct when we analyze the information acquisition stage.)

is no disclosure. Plainly, when s is disclosed, the DM will simply choose her most-preferred action, $\alpha(s|0) = \rho s$.

We start by fixing an arbitrary action $a \in \mathbb{R}$ the DM may choose in the event of nondisclosure, and ask whether the adviser will disclose his signal if he observes it, assuming that $B \geq 0$ (the logic is symmetric when $B < 0$). The answer can be obtained easily with the aid of Figure 1 below. The figure depicts, as a function of the signal, the action most preferred by the DM (ρs) and the action most preferred by the adviser ($\rho s + B$): each is a straight line, the latter shifted up from the former by the constant B. Since the DM will choose the action ρs whenever s is disclosed, the adviser will withhold s whenever the nondisclosure action a is closer to his most-preferred action, $\rho s + B$, than the disclosure action, ρs. This reasoning identifies the nondisclosure interval as the "flat" region of the solid line, which corresponds to the nondisclosure action chosen by the DM.

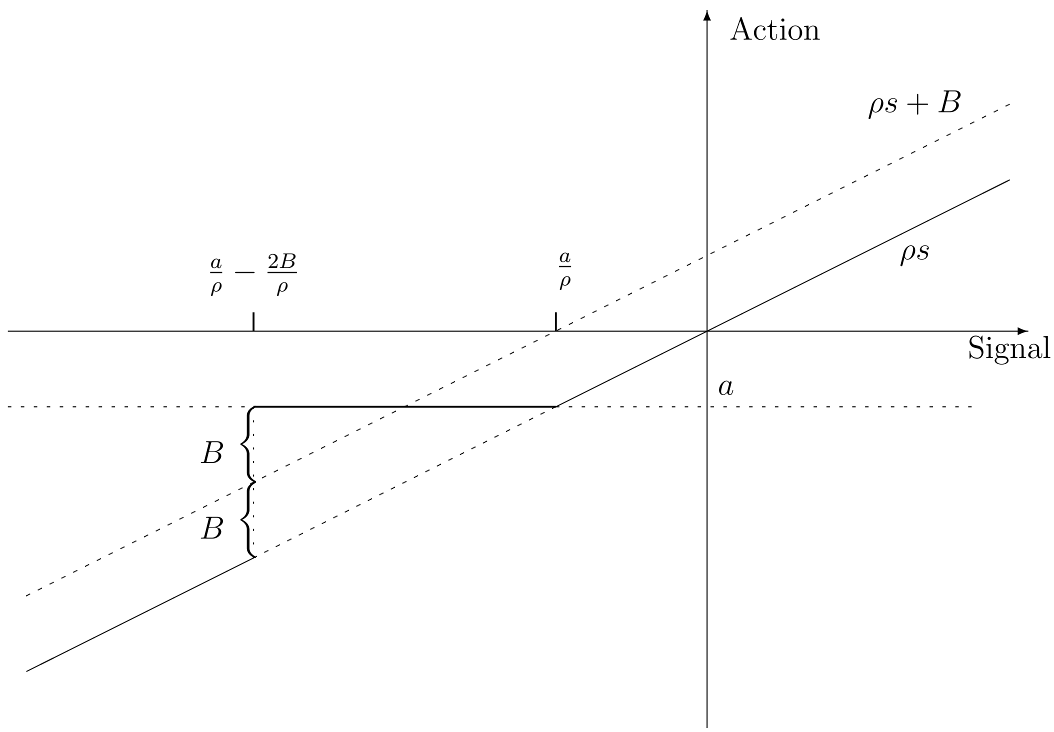

Figure 1: Optimal non-disclosure region

As seen in Figure 1, the adviser's best response is to withhold s (in case he observes s)

if and only if $s \in R(B,a) := [l\,(B,a)\,,h\,(a)]$ defined by

$$h\,(a) = \frac{a}{\rho}, \tag{2}$$

$$l(B,a) = h\,(a) - \frac{2B}{\rho}. \tag{3}$$

At $s = h(a)$, the DM will choose $a = \alpha\,(h(a)|0)$ whether s is disclosed or not, so the adviser is indifferent. At $s = l(B,a)$, the adviser is again indifferent between disclosure, which leads to $\alpha\,(l(B,a)|0) = a - 2B$, and nondisclosure, which leads to a, because they are equally distant from his most preferred action, $a - B$. For any $s \notin [l(B,a), h(a)]$, disclosure will lead to an action closer to the adviser's preferred action than would nondisclosure.[10]

Next, we characterize the DM's best response in terms of her nondisclosure action, for an arbitrary (measurable) set $S \subset \mathbb{R}$ of signals that the adviser may withhold. Her best response is to take the action that is equal to her posterior expectation of the state given nondisclosure, which is computed via Bayes rule:

$$a_N(p,S) = \frac{p\rho\int_S s\gamma\,(s;0)\,ds}{p\int_S \gamma\,(s;0)\,ds + 1 - p}, \tag{4}$$

where $\gamma(s;\mu)$ is a Normal density with mean μ and variance $\sigma_0^2 + \sigma_1^2$. Notice that the DM uses her own prior $\mu_{DM} = 0$ to update her belief. It is immediate that if S has zero expected value, then $a_N(p,S) = 0$. More importantly, for any $p > 0$, $a_N(p,S)$ increases when S gets larger in the strong set order.[11] Intuitively, the DM rationally raises her action when she suspects the adviser of not disclosing larger values of s.

An equilibrium of the disclosure sub-game requires that both the DM and the adviser must play best responses. This translates into a simple fixed point requirement:

$$S = R\,(B,a) \text{ and } a_N(p,S) = a. \tag{5}$$

Given any (B,p), let $(S(B,p), a_\emptyset(B,p))$ be a pair that satisfies (5), and let $\underline{s}(B,p)$ and $\overline{s}(B,p)$ respectively denote the smallest and the largest elements of $S(B,p)$. The following result ensures that these objects are uniquely defined; its proof, and all subsequent proofs not in the text, are in the Appendix.

[10] We assume nondisclosure when indifferent, but this is immaterial.

[11] A set S is larger than S' in the strong set order if for any $s \in S$ and $s' \in S'$, $\max\{s,s'\} \in S$ and $\min\{s,s'\} \in S'$.

Proposition 1. (Disclosure Equilibrium) *For any* (B, p)*, there is a unique equilibrium in the disclosure sub-game. In equilibrium, both* $\underline{s}(B, p)$ *and* $\overline{s}(B, p)$ *are equal to zero if* $B = 0$*, are strictly decreasing in* B *when* $p > 0$*, and strictly decreasing (increasing) in* p *if* $B > 0$ *(if* $B < 0$*). The nondisclosure action* $a_\emptyset(B, p)$ *is zero if* $B = 0$ *or* $p = 0$*, is strictly decreasing in* B *for* $p > 0$*, and is strictly decreasing (increasing) in* p *if* $B > 0$ *(if* $B < 0$*).*

It is straightforward that the adviser reveals his information fully to the DM if and only if $B = 0$, i.e. there is no interim bias. To see the effect of an increase in B (when $p > 0$), notice from (2) and (3) that if the DM's nondisclosure action did not change, the upper endpoint of the adviser's nondisclosure region would not change, but he would withhold more low signals. Consequently, by (4), the DM must adjust his nondisclosure action downward, which has the effect of pushing down both endpoints of the adviser's nondisclosure region. The new fixed point must therefore feature a smaller nondisclosure set (in the sense of strong set order) and a lower nondisclosure action from the DM. We call this the *prejudicial effect*, since a more upwardly biased adviser is in essence punished with a lower inference when he claims not to have observed a signal. The prejudicial effect implies in particular that for any $p > 0$ and $B \neq 0$, $a_\emptyset(B, p)B < 0$.

The impact of p can be traced similarly. An increase in p makes it more likely that nondisclosure from the adviser is due to withholding of information rather than a lack of signal. If $B > 0$ (resp. $B < 0$), this makes the DM put higher probability on the signal being low (resp. high), leading to a decrease (resp. increase) in the nondisclosure action, which decreases (resp. increases) the nondisclosure set in the strong set order.

Finally, it is worth emphasizing that the adviser's optimal disclosure behavior depends directly only on his interim bias, B, and the DM's nondisclosure action, $a_\emptyset$. In particular, it does not depend directly on the probability of acquiring a signal, although the the DM's belief about this probability affects the DM's nondisclosure action, and thereby, indirectly, the adviser's disclosure choice.

4 Opinions as Incentives

This section studies how the adviser's opinion affects his incentive to acquire information, and the implications this has on the optimal type of adviser for the DM. As a benchmark, the following Proposition establishes the fairly obvious point that, absent information acquisition concerns, the optimal adviser is a like-minded one.

Proposition 2. (Exogenous Effort) *If the probability of acquiring a signal is held fixed at some $p > 0$, the uniquely optimal type of adviser for the DM is like-minded, i.e. an adviser with $\mu = 0$.*

Proof. For any $p > 0$, $S(\mu, p)$ has positive measure when $\mu > 0$, whereas $S(0, p)$ has measure zero. Hence, the adviser $\mu = 0$ reveals the signal whenever she obtains one, whereas an adviser with $\mu \neq 0$ withholds the signal with positive probability. The result follows from the fact that DM is strictly better off under full disclosure than partial disclosure. ▌

We now turn to the the case where information acquisition is endogenous. To begin, suppose the DM believes that an adviser with type $\mu \geq 0$ will choose effort p^e. The following Lemma decomposes the payoff for the adviser from choosing effort p, denoted $U_A(p; p^e, B, \mu)$, in a useful manner.[12]

Lemma 1. *The adviser's expected utility from choosing effort p can be written as*

$$U_A(p; p^e, B, \mu) = K(B, \mu, p^e) + p\Delta(B, \mu, p^e) - c(p),$$

where

$$K(B, \mu, p^e) := -\int \left(a_\emptyset\left(B, p^e\right) - \left(\rho s + B\right)\right)^2 \gamma\left(s; \mu\right) ds - \tilde{\sigma}^2 \tag{6}$$

and

$$\Delta(B, \mu, p^e) := \int_{s \notin S(B, p^e)} \left[\left(a_\emptyset\left(B, p^e\right) - \left(\rho s + B\right)\right)^2 - B^2\right] \gamma\left(s; \mu\right) ds. \tag{7}$$

The first term in the decomposition, $K(\cdot)$, is the expected utility when a signal is not observed. Equation (6) expresses this utility by iterating expectations over each possible value of s, reflecting the fact that the DM takes decision $a_\emptyset(\cdot)$ without its disclosure whereas the adviser's preferred action if the signal were s is $\rho s + B$, and that $\tilde{\sigma}^2$ is the residual variance of the state given any signal. The second term in the decomposition, $p\Delta(\cdot)$, is the probability of a obtaining a signal multiplied by the expected gain from obtaining a signal. Equation (7) expresses the expected gain, $\Delta(\cdot)$, via iterated expectations over possible signals. To understand it, note that the adviser's gain is zero if a signal is not disclosed

[12] Even though the interim bias B is determined by μ, we write them as separate variables in the funcion $U_A(\cdot)$ to emphasize the two separate effects caused by changes in the difference of opinion: changes in prior beliefs over signal distributions and changes in the interim bias.

(whenever $s \in S(B, p^e)$), whereas when a signal is disclosed, the adviser's utility (gross of the residual variance) is $-B^2$, because the DM takes decision ρs.

We are now in a position to characterize the adviser's equilibrium effort level. Given the DM's belief, p^e, the adviser will choose p to maximize $U_A(p; p^e, B, \mu)$. By the Inada conditions on effort costs, this choice is in the interior of $[0, \bar{p}]$ and characterized by the first-order condition:

$$\frac{\partial U_A(p; p^e, B, \mu)}{\partial p} = \Delta(B, \mu, p^e) - c'(p) = 0.$$

Equilibrium requires that the DM's belief be correct, i.e. $p^e = p$. Therefore, in equilibrium, we must have

$$\Delta(B, \mu, p) = c'(p). \tag{8}$$

The following Lemma assures that (8) is necessary and sufficient for an equilibrium, and moreover, there is a solution to (8).

LEMMA 2. *For any (B, μ), there is a solution to (8), and p is an equilibrium effort choice if and only if $p \in (0, \bar{p})$ and satisfies (8).*

In general, we cannot rule out that there may be multiple equilibrium effort levels for a given type of adviser. The reason is that the DM's action in the event of nondisclosure depends on adviser's (expected) effort, and the adviser's equilibrium effort in turn depends on the DM's action upon nondisclosure.[13] For the remainder of the paper, for each (B, μ), we focus on the highest equilibrium effort. Since the interim bias B is uniquely determined by $B(\mu) = (1 - \rho)\mu$, we can define the equilibrium probability of information acquisition as a function solely of μ, which we denote by $p(\mu)$. Our first main result is:

PROPOSITION 3. (INCENTIVIZING EFFECT OF DIFFERENCE OF OPINION) *An adviser with a greater difference of opinion acquires information with higher probability: $p(\mu') > p(\mu)$ if $|\mu'| > |\mu|$.*

To see the intuition, first ignore the strategic disclosure of information, assuming instead that the outcome of the adviser's investigation is publicly observed. In this case, there is no

[13]Formally, multiplicity emerges when the function $\Delta(B, \mu, \cdot)$ crosses more than once with the strictly increasing function $c'(\cdot)$ over the domain $[0, 1]$. As we will discuss more shortly, if signals are public rather than privately observed by the adviser, there is a unique equilibrium because $\Delta(B, \mu, \cdot)$ is constant. Moreover, we show in the Appendix (in the proof of Proposition 3) that for all μ sufficiently close to 0, there is a unique equilibrium effort level.

prejudice associated with nondisclosure, so the DM will choose $a_\emptyset(B,p) = 0$ independent of B or p. It follows from a mean-variance decomposition that $-\sigma_0^2 - \mu^2$ is the expected utility for the adviser conditional on no signal, and $-\tilde{\sigma}^2 - (B(\mu))^2$ is the expected utility conditional on getting a signal. Hence, the adviser's marginal benefit of acquiring a signal, denoted $A^{pub}(\mu)$, is given by[14]

$$\Delta^{pub}(\mu) = \underbrace{\sigma_0^2 - \tilde{\sigma}^2}_{\text{uncertainty reduction}} + \underbrace{\mu^2 - (B(\mu))^2}_{\text{persuasion}}. \tag{9}$$

Acquiring information benefits the adviser by reducing uncertainty about the true state, as shown by the first part of (9). But in addition, the adviser expects to *persuade* the DM: without information, the adviser views the DM's decision as biased by μ, their ex-ante disagreement in beliefs; whereas with information, the disagreement is reduced to the interim bias, $B(\mu) = (1-\rho)\mu < \mu$. Since $\mu^2 - (B(\mu))^2$ is strictly increasing in $|\mu|$, the persuasion incentive is strictly larger for an adviser with a greater difference of opinion. This leads to such an adviser exerting more effort towards information acquisition. A related intuition is that the adviser expects action $\rho\mu$ conditional on acquiring information, whereas he knows that action 0 will be taken without information. Hence, the adviser believes that by acquiring information, he can persuade the DM to take an action that is closer in expectation to his own prior.[15] The benefit of such persuasion is more valuable to an adviser with greater difference of opinion.

Now consider the case where information is private, and the adviser strategically communicates. Suppose the DM expects effort p^e from the adviser of type μ. Then he will choose $a_\emptyset(B(\mu), p^e)$ when a signal is not disclosed. Since the adviser always has the option to disclose all signals, his marginal benefit of acquiring information and then strategically disclosing it, as defined by equation (7), is at least as large as the marginal benefit from (sub-optimally) disclosing all signals, which we shall denote $\Delta^{pri}(\mu, a_\emptyset(B(\mu), p^e))$. By

[14]Alternatively, one can also verify that equation (7) simplifies to equation (9) if the nondisclosure region $S(\cdot) = \emptyset$ and $a_\emptyset(\cdot) = 0$, as is effectively the case under public observation of signal.

[15]Of course, the DM does not expect to be persuaded: her expectation of her action conditional on a signal being acquired is 0. Instead, she expects that a signal will shift the adviser's preferred decision to shift towards her opinion. This feature that each player expects new information to persuade the other is also studied by Yildiz (2004) in a bargaining model with heterogeneous priors.

mean-variance decomposition again, we have

$$\begin{aligned}\Delta(B(\mu),\mu,p^e) &\geq \Delta^{pri}(\mu,a_\emptyset(B(\mu),p^e)) \\ &= \underbrace{\sigma_0^2-\tilde{\sigma}^2}_{\text{uncertainty reduction}} + \underbrace{\mu^2-(B(\mu))^2}_{\text{persuasion}} + \underbrace{(a_\emptyset)^2-2a_\emptyset\mu}_{\text{avoiding prejudice}}. \end{aligned} \tag{10}$$

Recall from Proposition 1 the prejudicial effect: for any $p^e > 0$ and $\mu \neq 0$, $a_\emptyset(B(\mu),p^e)\mu < 0$. Hence, for any $p^e > 0$ and $\mu \neq 0$, $\Delta^{pri}(\mu,p^e) > \Delta^{pub}(\mu)$: given that information is private, the DM's rational response to the adviser claiming a lack of information affects the adviser adversely—this is the prejudicial effect—and to avoid such an adverse inference, the adviser is even more motivated to acquire a signal than when information is public.

Propositions 1 and 3 identify the tradeoff faced by the DM: an adviser with a greater difference of opinion exerts more effort, but reveals less of any information he may acquire. Does the benefit from improved incentives for information acquisition outweigh the loss from strategic disclosure? We demonstrate below that this is indeed the case for at least some difference in opinion.

Proposition 4. (Optimality of Difference of Opinion) *There exists some $\mu_A \neq 0$ such that it is strictly better for the DM to appoint an adviser of type μ_A over a like-minded adviser.*

The optimality of difference of opinion is largely due to the persuasion effect. As the difference of opinion μ is raised slightly, the persuasion motive it generates creates a non-negligible benefit in increased information acquisition, whereas the prejudicial effect (which entails both communication loss and information acquisition gain) is negligible. This can be seen most clearly when the signal is perfectly informative, $\rho = 1$. In this case, $B(\mu) = 0$, so there is full disclosure in the communication stage, analogous to a situation where information is public; hence, appointing an adviser with difference of opinion is clearly desirable.[16] By continuity, there is a set of ρ's near 1 for which the adviser of type μ is better for the DM than the like-minded adviser. This argument verifies Proposition 4 for all ρ sufficiently close to 1. The proof in the Appendix shows that for any ρ, however far

[16]Formally, the DM's utility from appointing an adviser of type μ when $\rho = 1$ is

$$U_{DM}^{\rho=1}(\mu) := -\sigma_0^2(1-p(\mu)),$$

which is simply the ex-ante variance in the state multiplied by the probability of not acquiring a signal, because when a signal is acquired, there is no residual uncertainty. Proposition 3 implies that for any $\mu > 0$, $U_{DM}^{\rho=1}(\mu) > U_{DM}^{\rho=1}(0)$.

from 1, there is some adviser sufficiently near type 0 who is in fact better for the DM than an adviser of type 0.

Remark 1. *The conclusion of Proposition 4 does not depend on selecting the equilibrium with highest effort for a given adviser type. The proof of the Proposition establishes that for all μ sufficiently close to* 0*, there is a unique equilibrium effort level, and any adviser with $\mu \neq 0$ but sufficiently small in absolute value is strictly preferred to an adviser of type* 0*.*

Our main results are illustrated in Figures 2 and 3, which show numerical computations with the specified parameter values.[17] Figure 2 illustrates that the nondisclosure action is decreasing in μ, which happens for two reasons: directly through the prejudicial effect (holding p fixed, an increase in μ causes a decrease in $a_\emptyset$), and indirectly through the incentivizing effect, because an adviser with higher μ exerts more effort (when $\mu > 0$), leading to a higher p, which in turn also causes a decrease in $a_\emptyset$ (when $\mu > 0$) as noted in Proposition 1. Figure 3 illustrates that the DM never finds it optimal to appoint a like-minded expert. It also shows that as the signal gets more precise (so that information acquisition is more important to the DM), the DM may prefer to appoint an expert with greater difference of opinion.

5 Opinions and Preferences

We have thus far assumed that the DM and the available pool of advisers all have the same fundamental preferences, but differ in opinions. In this section, we augment the space of types to also allow for fundamental preference conflicts. This allows us to explore a number of issues, such as: Will the DM benefit from an adviser with different preferences in the same way she will benefit from one with a different opinion? If an adviser can be chosen from a very rich pool of advisers differing both in opinions and preferences, how will the DM combine the two attributes? For instance, for an adviser with a given preference, will she prefer him to be a *skeptic*—one who doubts that discovering the state of the world will shift the DM's action in the direction of his preferences bias—or a *zealot*—one who believes that his preference will also be "vindicated by the evidence."

To keep matters simple, suppose, as is standard in the literature, that a player's preferences are indexed by a single bias parameter $b \in [\underline{b}, \overline{b}]$, with $\underline{b} < 0 < \overline{b}$, such that his

[17] The figures only show positive values of μ, since there is an obvious symmetry for negative values.

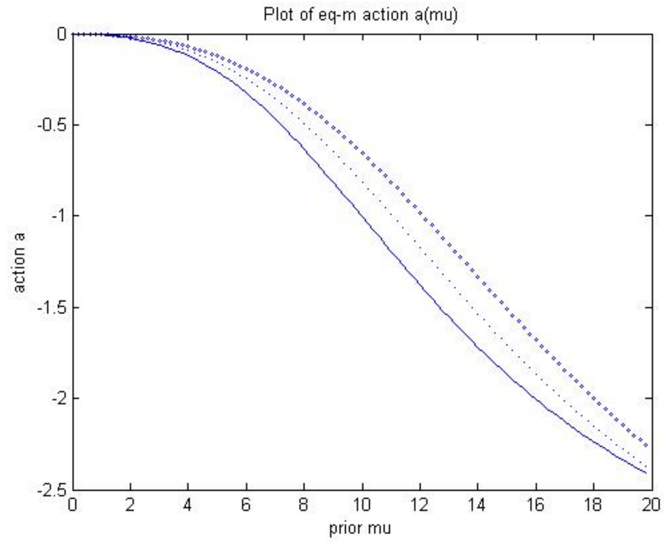

Figure 2: Equilibrium non-disclosure action as a function of adviser's opinion. Parameters: $c(p) = \frac{p^4}{(1-p)^2}$; $\sigma_1^2 = 2$; highest curve corresponds to $\rho = \frac{9}{10}$, middle curve to $\rho = \frac{8}{9}$, and lowest curve to $\rho = \frac{7}{8}$.

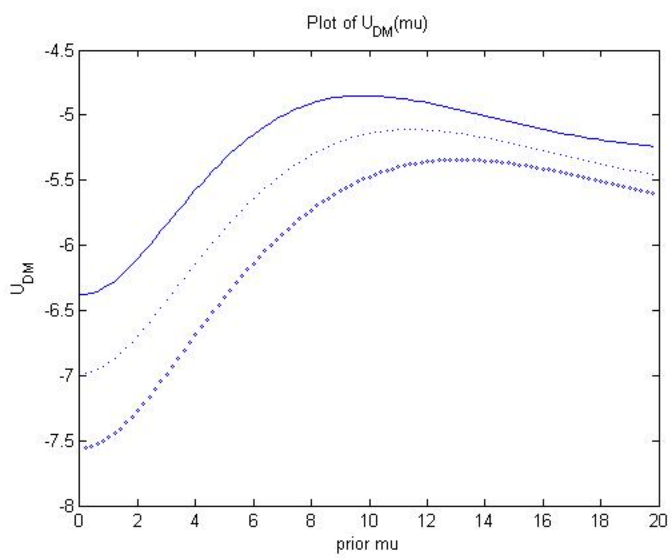

Figure 3: DM's utility as a function of adviser's opinion. Parameters: $c(p) = \frac{p^4}{(1-p)^2}$; $\sigma_1^2 = 2$; highest curve corresponds to $\rho = \frac{7}{8}$, middle curve to $\rho = \frac{8}{9}$, and lowest curve to $\rho = \frac{9}{10}$.

state-dependent von Neumann-Morgenstern utility is $u(a, \omega, b) = -(a - \omega - b)^2$. The adviser therefore now has a two-dimensional type (that is common knowledge), (b, μ). The DM's type is normalized as $(0, 0)$.

Interim but not ex-ante equivalence. Similar to earlier analysis, it is straightforward that an adviser of type (b, μ) desires the action $\alpha(s|b, \mu) := \rho s + (1 - \rho)\mu + b$ when signal s is observed. Hence, such an adviser has an interim bias of $B(b, \mu) := (1 - \rho)\mu + b$. This immediately suggests the interchangeability of the two kinds of biases—preferences and opinions—in the disclosure sub-game. For any adviser with opinion bias μ and no preference bias, there exists an adviser with only preference bias $b = (1-\rho)\mu$ such that the latter will have precisely the same incentives to disclose the signal as the former. Formally, given the same effort level, the disclosure sub-game equilibrium played by the DM and either adviser is the same.

This isomorphism does not extend to the information acquisition stage. To see this, start with an adviser of type $(0, \mu)$, i.e., with opinion bias μ but no preference bias. When such an adviser does not acquire a signal, he expects the DM to make a decision that is distorted by at least μ from what he regards as the right decision.[18] Consider now an adviser of type $(\mu, 0)$, i.e., with preference bias $b = \mu$ and no opinion bias. This adviser also believes that, absent disclosure of a signal, the DM will choose an action that is at least μ away from his most preferred decision. Crucially, however, their expected payoffs from disclosing a signal are quite different. The former type (opinion-biased adviser) believes that the signal will vindicate his prior and thus bring the DM closer toward his ex-ante preferred decision; whereas the latter type (preference-biased adviser) has no such expectation. One concludes that *the persuasion motive does not exist for an adviser biased in preferences alone.*

Publicly observed signal. To see how the two types of biases can interact in affecting the incentive for information acquisition, it is useful to first consider the case where the adviser's signal (or lack thereof) is publicly observed. This makes matters simple, because there is no strategic withholding of information. Fix any adviser of type (b, μ). If no signal is observed, the DM takes action 0, while the the adviser prefers the action $b+\mu$. Hence, the adviser has expected utility $-\sigma_0^2 - (b+\mu)^2$. If signal s is observed, then the DM takes action ρs; since the adviser prefer action $\rho s + B(b, \mu)$, he has expected utility $-\tilde{\sigma}^2 - (B(b, \mu))^2$.

[18] "At least," because the prejudicial effect will cause the DM to take an action even lower than 0, unless information is public or signals are perfectly-informative.

The adviser's expected gain from acquiring information is, therefore,

$$\Delta^{pub}(b,\mu) = \underbrace{\sigma_0^2 - \tilde{\sigma}^2}_{\text{uncertainty reduction}} + \underbrace{\left(2\rho - \rho^2\right)\mu^2}_{\text{persuasion}} + \underbrace{(1+\rho)\,b\mu.}_{\text{reinforcement}} \tag{11}$$

Suppose first $\mu = 0$, so the adviser is like-minded. In this case, $\Delta^{pub}(b,0)$ is independent of b. That is, the incentive for a like-minded adviser to acquire information does not depend on his preference, and consequently, there is no benefit to appointing an adviser who differs only in preference. This stands in stark contrast to the case of difference of opinion, $(0,\mu)$, $\mu \neq 0$, where equation (9) showed that advisers with greater difference of opinion have bigger marginal benefits of acquiring information, and are therefore strictly better for the DM under public information. This clearly shows the distinction between preferences and opinions.

Now suppose $\mu \neq 0$. Then, the persuasion effect reappears, as is captured by the second term of (11). More interestingly, the adviser's preference also matters now, and in fact interacts with the opinion bias. Specifically, *a positive opinion bias is reinforced by a positive preference bias, whereas it is counteracted by a negative preference bias*; this effect manifests itself as the third term in (11). The intuition turns on the concavity of the adviser's payoff function, and can be seen as follows. Without a signal, the adviser's optimal action is away from the DM's action by $|b+\mu|$. Concavity implies that the bigger is $|b+\mu|$, the greater the utility gain for the adviser when he expects to move the DM's action in the direction of his ex-ante bias. Therefore, when $\mu > 0$, say, an adviser with $b > 0$ has a greater incentive to acquire information than an adviser with $b < 0$. In fact, if b were sufficiently negative relative to $\mu > 0$, the adviser may not want to acquire information at all, because he expects it to shift the DM's decision *away* from his net bias of $b + \mu$.

Privately observed signal. When the signal is observed privately by the adviser, the prejudicial motive is added to the adviser's incentive for information acquisition. The next proposition states an incentivizing effect of both preference and opinion biases. Extending our previous notation, we use $p(B,\mu)$ to denote the highest equilibrium effort choice of an adviser with interim bias B and prior μ.

Proposition 5. *Suppose* $(|B(b,\mu)|, |\mu|) < (|B(b',\mu')|, |\mu'|)$ *and* $B(b',\mu')\mu' \geq 0$.[19] *Then,* $p(B(b',\mu'),\mu') > p(B(b,\mu),\mu)$.

[19] We follow the convention that $(x,y) < (x',y')$ if $x \leq x'$ and $y \leq y'$, with at least one strict inequality.

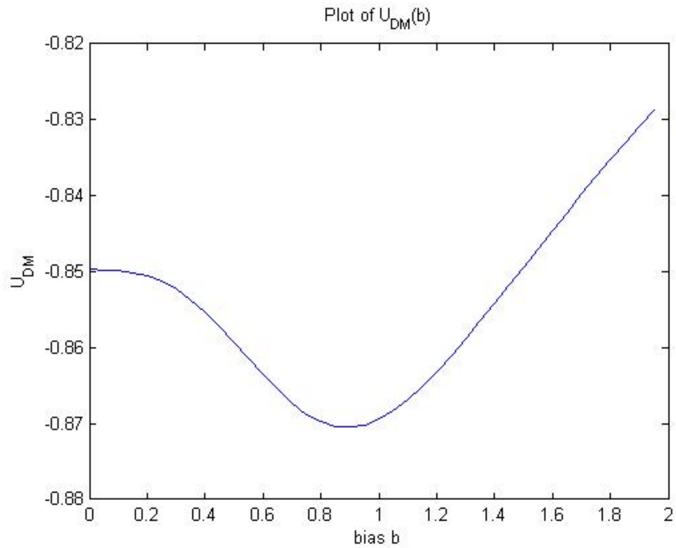

Figure 4: DM's utility as a function of adviser's preference. Parameters: $c(p) = \frac{p^2}{1-p}$, $\sigma_1^2 = 1$, $\sigma_0^2 = 0.5$.

Proposition 5 nests Proposition 3 as a special case with $b = b' = 0$. Setting $\mu = \mu' = 0$ gives the other special case in which the adviser differs from the DM only in preference. Unlike under public information, a preference bias alone creates incentives for information acquisition when the outcome of the adviser's experiment is private. The reason is that an adviser exerts additional effort to avoid the prejudicial inference the DM attaches to nondisclosure. Of course, from the DM's point of view, this incentive benefit is offset by the loss associated with strategic withholding of information. It turns out that these opposing effects are of the same magnitude locally when $|b| \approx 0$. Hence, a difference of preference is not unambiguously beneficial to DM in the way that difference of opinion is. Indeed, a numerical example shows that the DM's utility is decreasing in $|b|$ around $b = 0$, but interestingly, starts increasing when $|b|$ becomes sufficiently large, to the point where it can rise above the utility associated with type $b = 0$. This is shown in Figure 4. In such cases, the DM never prefers an adviser with preference bias unless the bias is sufficiently large, contrasting with difference of opinion. This difference may matter if the space of available adviser types is not sufficiently large (such as $\bar{b} < 1.4$ in the example plotted in Figure 4).

More generally, Proposition 5 shows how the two types of biases interact with respect to the incentive for information acquisition. The following corollaries record the nature of the interaction.

Corollary 1. (Complementarity of opinion and preference) *If* $(b', \mu') > (b, \mu) \geq$

0, *then an adviser with* (b', μ') *choose a higher effort than one with* (b, μ)*.*

Thus, in the domain $(b, \mu) \in \mathbb{R}^2_+$, an increase in either kind of bias—preference or opinion—leads to greater information acquisition.

Corollary 2. (Zealot vs. skeptic) *Suppose an adviser has type* (b, μ) *such that* $B(b, \mu) \geq 0$ *but that* $\mu < 0$*. Replacing the adviser with one of type* $(b, -\mu)$ *leads to a higher effort.*

An adviser of type (b, μ) with $B(b, \mu) \geq 0$ but $\mu < 0$ likes actions higher than the DM would like if the state of the world were publicly known, yet he is a priori pessimistic about obtaining a signal that will shift the DM's action upward. In this sense, he is a *skeptic*, and does not have a strong incentive for information acquisition. Replacing him with a *zealot* who believes that information about the state will in fact lead the DM to take a higher action leads to more information acquisition.

The final corollary shows that having access to a rich pool of advisers on both opinion and preference dimensions endows the DM with enough degree of freedom to eliminate disclosure loss altogether, and yet use the adviser's type as an incentive instrument.

Corollary 3. (Optimal type) *If* $B(b, \mu) = B(b', \mu') \geq 0$ *and* $\mu' > \mu \geq 0$*, then the adviser with* (b', μ') *chooses a higher effort than the one with* (b, μ)*. Moreover, the DM strictly prefers appointing the former. In particular, if one raises* μ *and lowers* b *so as to maintain* $B(b, \mu) = 0$*, then an higher effort is induced while maintaining full disclosure.*

Choosing an adviser who has opinion $\mu > 0$ but negative preference bias $b = -(1 - \rho)\mu$ eliminates interim bias altogether, and thus avoids any strategic withholding of information. If this can be done without any constraints, the DM can raise μ unboundedly and increase his expected utility. However, since the space of types is likely bounded (as is assumed), it may be optimal for the DM to choose an expert with an interim bias $B(b, \mu) > 0$, as was the case when advisers are differentiated by opinions alone.

6 On Delegation and Confidence

In this section, we discuss two other issues that can be raised in our framework. For simplicity, we return to the baseline setting where individuals are only distinguished by their opinions, sharing the same fundamental preferences.

6.1 Delegation

An important issue in various organizations is the choice between delegation and communication.[20] One prominent view is that delegation of decision-making authority to an agent increases his incentives to become informed about the decision problem if he can extract a greater fraction of the surplus from information gathering, but is costly to the DM insofar as the agent's actions when given authority do not maximize the DM's interests. Aghion and Tirole (1997) label this the "initiative" versus "loss of control" tradeoff.

Delegation involves a costly loss of control in our model as well when the adviser has a different opinion. This is because the decision taken by such an adviser will not coincide with the DM's preferred action, regardless of whether a signal is observed or not. But in addition, delegation also fails to harness the incentive benefits associated with difference of opinion under communication. The reason is that the adviser, in possession of decision-making authority, has no incentive to acquire information to either persuade the DM or to avoid prejudice. Regardless of his opinion, an adviser would acquire information solely to reduce uncertainty about the state. This implies that if the DM must delegate, she would strictly prefer to choose a like-minded adviser, thereby avoiding costly loss of control. Notice, however, that delegation and communication are equivalent when the adviser is like-minded (since there is no ex-ante nor interim bias). Our earlier results shows that the DM would strictly prefer to communicate, instead, with an appropriately chosen adviser. The following result summarizes this discussion:

PROPOSITION 6. *Under delegation, it is uniquely optimal for the DM to choose a like-minded adviser. However, such an arrangement is strictly worse for the DM than retaining authority and choosing an appropriate adviser with a difference of opinion.*

6.2 Heterogeneity in Confidence

Suppose that rather than being differentiated by prior mean, all advisers have the same prior mean as the DM ($\mu = 0$), but they differ in confidence in their ability, represented by beliefs about their signal precision. It is convenient to map signal precision into the weight

[20]Our interest here concerns situations where there is an information acquisition problem, but Dessein (2002) shows that the question is subtle even when information is exogenous. Also, we only consider the two extremes: full delegation and no delegation, ignoring constrained delegation. A standard justification is that of incomplete contracting. Szalay (2005) analyzes a related mechanism design problem with commitment, but our setting differs because of verifiable information and because differences of opinion induce interim bias.

the posterior mean places on signal versus prior mean, as before. That is, if an adviser believes his signal has variance σ_1^2, we represent him via $\rho(\sigma_1^2) := \frac{\sigma_0^2}{\sigma_0^2+\sigma_1^2}$.[21] The DM has $\rho_{DM} \in (0,1)$, and she can choose an adviser with any $\rho_A \in [0,1]$. Interpreting ρ_{DM} as a baseline, an adviser with $\rho_A > \rho_{DM}$ is overconfident of his ability and an adviser with $\rho_A < \rho_{DM}$ is underconfident.[22]

The preferred action for an adviser of type ρ_A who observes signal s is $\rho_A s$, whereas the DM would take action $\rho_{DM} s$ if the signal is disclosed. Hence, any adviser with $\rho_A \neq \rho_{DM}$ has a conflict of interest in the disclosure sub-game. Surprisingly however, so long as $\rho_A \geq \rho_{DM}$, there is an equilibrium in the disclosure sub-game that fully reveals the outcome of the adviser's investigation, for any effort choice. To see this, notice that if the DM believes the adviser never withholds his signal, he optimally plays $a_\emptyset = 0$. It is then optimal for the adviser to disclose any signal he acquires because $|\rho_A s| \geq |\rho_{DM} s|$ for all s (with strict inequality when $s \neq 0$ and $\rho_A > \rho_{DM}$).[23]

We can show that the incentives to acquire information are strictly higher the more overconfident an adviser is. The intuition is that he believes the value of a signal is larger, and hence perceives a larger marginal benefit of effort, even after accounting for the fact that the DM's decision does not respond to a signal as much as he would like. Therefore, among overconfident advisers, the DM strictly prefers one with $\rho_A = 1$. In the proof of the following result, we establish that an underconfident adviser is never optimal either.

Proposition 7. *If advisers are distinguished only by confidence, ρ_A, the DM uniquely prefers to appoint a maximally overconfident adviser, i.e. one with $\rho_A = 1$ who believes that his signal is perfectly informative.*

7 Related Literature

Formally, our model builds upon a Sender-Receiver costless signaling game of verifiable information, first studied by Grossman (1981) and Milgrom (1981). The "unraveling"

[21] This makes it clear that the subsequent analysis can also be applied to differences in beliefs about the variance of the prior distribution of state of the world.

[22] Admati and Pfleiderer (2004) and Kawamura (2007) analyze cheap talk games with an over/underconfident Sender. Information acquisition is not endogenous in either paper.

[23] If there is any other equilibrium in the disclosure sub-game, we assume the fully revealing one is played. Aside from being intuitively focal and simple, it is the only disclosure equilibrium that is "renegotiation proof" at the point where effort has been exerted but investigation outcome not yet observed (regardless of how much effort has/is believed to have been exerted).

phenomenon noted by these authors does not arise here because the Sender (adviser) may not possess information, as pointed out by Shin (1994, 1998), among others.

The idea that differences of opinion induce interim conflicts, in turn entailing communication loss, is not surprising. In most models of strategic communication with exogenous information, if the DM were permitted to choose the type of adviser she communicates with, she would choose an adviser who shares her interim preferences (Suen, 2004, explictly discusses this point).[24] Our point of departure is to endogenize the acquisition of information. Early papers that incorporate this aspect include Matthews and Postlewaite (1985) and Shavell (1994); they do not focus on the tradeoff between information acquisition and transmission, and how this is affected by the degree of conflict.

We now relate our findings to a few recent papers that have explored related and complementary themes to ours. Dur and Swank (2005) show that when faced with a *binary decision*, an open-minded or moderate adviser values information more than an ex-ante biased adviser. Thus, an ex-ante biased DM may choose a more moderate adviser to encourage information acquistion, even though this adviser will engage in strategic communication. However, an ex-ante unbiased DM will find it optimal to choose an adviser of the same type. In contrast, in our model, the decision space is continuous and unbounded, hence there is no notion of being moderate or open-minded; instead what the DM benefits from is a difference of opinion, and this is beneficial regardless of the DM's type.

Gerardi and Yariv (2007a) show, in a jury model with binary decisions and *public information*, that appointing a juror with a different preference (in terms of the threshold of reasonable doubt) from the DM can be beneficial, and in fact, the optimal juror is typically one who is extremely biased in the opposite direction of the DM.[25] The intuition is simple: a juror who is ex-ante biased in the opposite direction of the DM faces an unfavorable status quo. Hence, he can only gain from new information, which will at worst not change the status quo or, with luck, shift the DM's decision to his favored alternative.[26] With non-binary decisions, however, new information can lead to a decision that is less favorable than the status quo. In fact, without a restriction on the decision space, new information does not

[24] This is also true in models with multiple advisers, such as Krishna and Morgan (2001) and Li and Suen (2004).

[25] See also Section 5.2 of Dur and Swank (2005).

[26] A related point is made by Hori (2006) who considers an agent hired to advise a principal on the adoption of a new project. The cost of providing incentives for information acquisition may be lower for an agent who is biased toward adoption, if the project is a priori unlikely to be profitable (in which case an unbiased agent may simply recommend non-adoption without learning about the project).

shift the decision on average when the conflict is one of preferences alone. Consequently, as was shown in Section 5, a preference conflict alone with public information cannot motivate information acquisition when the decision is continuous. Gerardi and Yariv (2007a) discuss the appointment of multiple jurors and sequential consultation.

The persuasion effect that we find from difference of opinion is also noted by Van den Steen (2004).[27] His model differs from ours, however, in a number of ways; for example he assumes both binary signals and decisions. Substantively, he does not study the prejudicial effect that occurs when information is privately observed by the adviser, and its implications for information acquisition. There is, therefore, no analog of our finding that some difference of opinion is optimal for the DM even accounting for the cost of strategic disclosure. Instead, Van den Steen (2004) studies coordination issues between multiple individuals' action choices.

The optimality of having an adviser whose interim preferences are different from the DM's is reminiscent of and broadly related to Dewatripont and Tirole (1999) and Prendergast (2007), but the frameworks and forces at work are quite different. Dewatripont and Tirole (1999) study the optimality of giving monetary rewards for "advocacy" from agents when the central problem is that of multitasking between conflicting tasks. In contrast, the effects in our model are derived solely from a single effort choice. Prendergast (2007) shows that a society may prefer to appoint bureaucrats with preferences different from its own, because such agents may have greater "intrinsic motivation" to exert effort. For example, in a setting with no conflicts of interest at all, society would appoint bureaucrats who are more altruistic towards their clients than the average member of society, simply because such bureaucrats care more about making good decisions and thus will exert more effort.

Finally, the current paper is consistent with the emerging theme that ex-post suboptimal mechanisms are sometimes ex-ante optimal because they can provide greater incentives for agents to acquire information. Examples in settings related to ours are Li (2001) and Szalay (2005).[28] Although we are interested in environments where contracting is essentially infeasible, our results can be interpreted in this light. It is important to stress, however, that *ex post* suboptimality is not the only reason that it is *ex ante* optimal for the DM to appoint an adviser with a difference of opinion: even if signals are public when acquired—in

[27] This was brought to our attention after the first draft of our paper was complete; we thank Eric Van den Steen.

[28] Examples in other environments include Bergemann and Valimaki (2002), Gerardi and Yariv (2007b), and Gershkov and Szentes (2004).

which case the DM always takes an optimal decision given the adviser's information—an adviser who has a larger difference of opinion with the DM has a greater incentive to acquire information, because of the persuasion motive.

8 Concluding Remarks

This paper has developed a model of strategic communication to study the costs and benefits of difference of opinion from an incentive point of view. Our main findings are threefold. First, because difference of opinion leads to interim conflict of interest even when fundamental preferences agree, it leads to a loss in information transmission via strategic disclosure. This leads to a prejudicial effect against an adviser who has a difference of opinion unless he provides information. Second, difference of opinion increases the incentives for an adviser to exert costly effort towards information acquisition. This is for two reasons: a motive to avoid the prejudicial inference associated with nondisclosure, and a belief that he can systematically shift the decision maker's action in the direction of his opinion by acquiring information (the persuasion motive). Third, the decision maker optimally resolves the tradeoff by appointing an adviser who has at least some difference of opinion with her.

We now mention a few applications that our analysis provides some insight into.

Independent counsel law. The US Ethics in Government Act, enacted in 1978 amid the Watergate scandal, created the office of independent counsel to investigate allegations of misconduct by high-ranking government officials. While the office was given independence, broad jurisdiction and unlimited budget to insure its impartiality, it was precisely its neutrality, along with its investigation tactics and (over-)spending, that has been often questioned. The office, vilified by some as the "fourth branch" of the government, was expired along with the law in 1999.

One can think of a counsel as the adviser in our model, who reports to Congress and/or public opinion—the DM in our model—who in turn forms its final opinion or metes out appropriate punishment (the action in our model). Suppose a right-leaning counsel is appointed to investigate a Democrat president. When such a counsel fails to turn up incriminating evidence against the president, public opinion will be rather favorable to the president, reflecting a prejudicial effect: the rational suspicion that such a counsel may have concealed exculpatory information. As we argued, the desire to avoid such a prejudicial inference can motivate the counsel to engage in excessive effort. Any impression of such excess (high p in our model) in turn shifts the public's inference further against the counsel and in favor of the president. This observation is consistent with the shifting of the scrutiny from Bill Clinton to the independent counsel Kenneth Starr in the wake of the latter's investigation of Whitewater and related matters, and the politicization of the investigation.[29]

Maverick sciences. History of sciences is rich with episodes in which "maverick sciences" or unorthodox views are dismissed by the majority of the scientific community and/or general public for a while, but eventually become accepted as a part of mainstream science. Galileo's struggle for heliocentricism against the church and community at large is particularly famous, as mentioned in our introduction, but there are many other well-known cases.[30] The role of maverick science in furthering scientific knowledge would be consistent with our analysis. Interpreting the public or general scientific community as the DM in our model, one can ask how well-served she is by a particular scientist (adviser), even if she does not volitionally choose him. Our results imply that new discoveries are more likely to come from those with different views about the world (than would be explainable by their population composition), and that public interest can be served by such individuals. Moreover, our theory also suggests that establishments are right to dismiss non-conformist theories until rigorously proven, both because scientists with unorthodox views may not reveal evidence that is contrary to their views, and because it is often the need to persuade and process of persuasion itself that leads to scientific breakthroughs.

Leader with different opinions. We have discussed difference of opinion as a desirable quality when selecting an adviser. But, a similar point applies when selecting a leader or a representative of a group. Suppose an organization has agents of given type but may be looking for a leader, the DM in our model. This is the task corporate boards are often faced with when finding a successor to a CEO who either retires or steps down. Often an important issue is whether to find somebody, perhaps from inside, who conforms to views of the existing managers or somebody from outside who does not share the same

[29] In a recent show of support for Hilary Clinton, Bill Clinton said "Ken Starr spent \$70 million and indicted innocent people to find out that I wouldn't take a nickel to see the cow jump over the moon." ("Campaigning for His Wife, Shadowed by Past Battles" in the *Washington Post*, January 18, 2008.)

[30] For example, the germ theory of disease, the theory of continental drift, and black holes were all held in suspicion for long periods of time before they became commonly accepted.

vision about the future direction of the company and thus will bring a radical change to the corporation. Our results can shed some light on such a choice. They imply that, even though a board may not be convinced about the need for change, they may nevertheless decide to bring a leadership with a different vision, so as to induce the firms' agents to exert more effort in searching for a better direction for the future of the company. Of course, based on our discussion of delegation, this benefit must be traded off against the cost of the leader implementing decisions the board may not agree with.

Finally, we conclude by briefly discussing some extensions.

The nature of effort. We have assumed that effort is success-enhancing in the sense of increasing the probability of observing an informative signal about the state. Alternatively, one could assume that there is a fixed probability with which the adviser observes a signal, but higher effort increases the precision of the observed signal. We believe our insights would extend to this setting. The incentivizing effect of difference of opinion and non-optimality of a like-minded expert certainly hold when the signal is publicly observed, because in this case there is no disclosure loss, and an adviser with a more extreme opinion has more to gain by producing a more precise signal.

Soft information. We have treated information as hard or verifiable. Aside from the technical tractability this offers, the assumption seems appropriate in various applications. On the other hand, there are some situations where information may be soft or unverifiable. Although a full analysis is left to future research, we can offer some observations. It is important to distinguish between two possibilities here.

First, as in Austen-Smith (1994), suppose that the adviser can prove that he has acquired a signal, but statements about what the signal is are cheap talk. Our main insights appear to extend to this setting, subject to the caveat of multiple equilibria in the communication stage owing to the usual difficulties with cheap talk. In particular, when signals are perfectly revealing of the state, there is an equilibrium in the communication game where the adviser (regardless of his opinion) reveals whether he is informed or not, and if informed, perfectly reveals his signal. Loosely, this is the "most informative" communication equilibrium. Given its selection, a greater difference of opinion leads to greater information acquisition, because of the persuasion motive, while there is no information loss from communication. By continuity, some difference of opinion will be optimal when signals are close to perfect, even though there is a loss of information from strategic communication in this case.

On the other hand, if the adviser cannot even prove that he has acquired a signal, so that all statements are cheap talk, the incentivizing effects of differences of opinion may be significantly mitigated. Intuitively, the persuasion motive from difference of opinion creates incentives because the adviser believes that by exerting more effort, he creates the ability to systematically shift the DM's action in his desired direction. Under complete cheap talk, effort does not affect the set of claims the adviser can make, casting doubt on whether the incentivizing effect exists.

Multiple advisers. We have studied a setting where a single adviser must be chosen. Decision makers sometimes consult more than one adviser, for instance, by forming a committee of advisers. Multiple advisers raise new issues both in terms of possible free riding amongst the advisers and the DM's response in committee design. We conjecture that the main thrust of our paper—the optimality of appointing advisers who differ in opinions from the DM—would extend to such a setting. A particularly interesting question is whether the optimal committee should consist of congruent or opposing agents, i.e. agents whose opinions are on the same side relative to the DM, or on opposite sides. We hope to address this in subsequent work.

A Appendix: Proofs

Proof of Proposition 1: By assumption, $p < 1$. Without loss, we assume $B \geq 0$. A symmetric argument will establish the result for the opposite case of $B < 0$. We start by deriving an equation whose solution will constitute the equilibrium condition (5). First, it follows from (5) that

$$a_{\emptyset}(B,p) = a_N(p, [\underline{s}(B,p), \overline{s}(B,p)]).$$

Substituting in from (2), (3), and (4) gives the main equation:

$$\overline{s}(B,p) = \frac{p}{p\int_{\overline{s}(B,p)-\frac{2B}{\rho}}^{\overline{s}(B,p)} \gamma(s;0)\,ds + 1 - p} \int_{\overline{s}(B,p)-\frac{2B}{\rho}}^{\overline{s}(B,p)} s\gamma(s;0)\,ds. \tag{12}$$

We will show that there is a unique solution to (12) in two steps below; this implies that is a unique disclosure equilibrium.

Step 1. *For any p, $\overline{s}(0,p) = \underline{s}(0,p) = a_{\emptyset}(0,p) = 0$.*

Proof: Immediate from the observation that $l(0,a) = h(a)$, and $a_N(p,S) = 0$ if S has measure 0. ‖

Step 2. *For any (B,p), there is a unique equilibrium in the disclosure game.*

Proof: Step 1 proves the result for $B = 0$, so we need only that show that there is a unique solution to (12) when $B > 0$. This latter is accomplished by showing that there is a unique solution to

$$\Upsilon(\overline{s}; B, p) := -p\int_{\overline{s}-\frac{2B}{\rho}}^{\overline{s}} (\overline{s} - s)\gamma(s;0)\,ds - \overline{s}(1-p) = 0. \tag{13}$$

Without loss of generality, we can restrict attention to $\overline{s} < 0$, because there is no solution to (13) with $\overline{s} \geq 0$ when $B > 0$. To see that there is at least one solution, apply the intermediate value theorem with the following observations: $\Upsilon(\overline{s}; B, p)$ is continuous in $\overline{s}$, and satisfies $\Upsilon(0; B, p) < 0$ and $\Upsilon(\overline{s}; B, p) \to \infty$ as $\overline{s} \to -\infty$ (because the integral in (13) is positive and bounded above by $4(\frac{B}{\rho})^2$, and $p < 1$).

To prove uniqueness, observe

$$\begin{aligned}\frac{\partial}{\partial \overline{s}}\Upsilon\left(\overline{s};B,p\right) &= p\left(1+2(B/\rho)\gamma\left(\overline{s}-2(B/\rho);0\right)-\int_{\overline{s}-2(B/\rho)}^{\overline{s}}\gamma\left(s;0\right)ds\right)-1 \\ &< p\left(1+2(B/\rho)\gamma\left(\overline{s}-2(B/\rho);0\right)-\int_{\overline{s}-2(B/\rho)}^{\overline{s}}\gamma\left(\overline{s}-2(B/\rho);0\right)ds\right)-1 \\ &= p-1\leq 0,\end{aligned}$$

where the inequality uses the fact that $\gamma\left(\cdot;0\right)$ is strictly increasing on the negative Reals. Consequently, there can only be one solution to (13). $\|$

The comparative statics results are established in several steps again for the case $B \geq 0$ (with the symmetric argument applicable for the opposite case $B < 0$).

Step 3. *For any* $(B,p) >> (0,0)$, $\frac{\partial}{\partial B}\overline{s}(B,p) < 0$.

Proof: We showed earlier that $\frac{\partial}{\partial \overline{s}}\Upsilon\left(\overline{s};B,p\right) < 0$. We also have

$$\frac{\partial}{\partial B}\Upsilon\left(\overline{s},B,p\right) = -\frac{4Bp}{\rho^2}\gamma\left(\overline{s}-2(B/\rho);0\right) < 0.$$

By the implicit function theorem,

$$\frac{\partial \overline{s}\left(B,p\right)}{\partial B} = -\frac{\frac{\partial \Upsilon\left(\overline{s}(B,p);B,p\right)}{\partial B}}{\frac{\partial \Upsilon\left(\overline{s}(B,p);B,p\right)}{\partial \overline{s}}} < (=)0,$$

(if $B=0$). $\|$

Step 4. *For any* $B \geq 0, p > 0$, $\frac{\partial}{\partial B}\underline{s}(B,p) < 0$.

Proof: The result follows from Step 3, upon noting that $\underline{s}(B,p) = \overline{s}(B,p) - \frac{2B}{\rho}$. $\|$

Step 5. *For any* $p > 0$, $\frac{\partial}{\partial p}\overline{s}(B,p) < (=)\, 0$ *if* $B > (=)\, 0$*; and* $\frac{\partial}{\partial p}\underline{s}(B,p) < (=)\, 0$ *if* $B > (=)\, 0$.

Proof: We showed earlier that $\frac{\partial}{\partial \overline{s}}\Upsilon\left(\overline{s};B,p\right) < 0$. We also have

$$\frac{\partial}{\partial p}\Upsilon\left(\overline{s};B,p\right) = -\int_{\overline{s}-2(B/\rho)}^{\overline{s}}\left(\overline{s}-s\right)\gamma\left(s;0\right)ds + \overline{s}.$$

By the implicit function theorem, $\frac{\partial \overline{s}(B,p)}{\partial p} = -\frac{\frac{\partial \Upsilon(\overline{s}(B,p);B,p)}{\partial p}}{\frac{\partial \Upsilon(\overline{s}(B,p);B,p)}{\partial \overline{s}}}$. The first statement is proven by noting that $\overline{s}(0,p) = 0$ and that, for any $B > 0$, $\overline{s}(B,p) \leq 0$. The second statement follows from the first statement, since $\underline{s}(B,p) = \overline{s}(B,p) - \frac{2B}{\rho}$. $\|$

Step 6. *The nondisclosure action $a_\emptyset(B,p)$ is zero if $B=0$ or $p=0$, is strictly decreasing in B for $p>0$, and is strictly decreasing (increasing) in p if $B>0$ (if $B<0$).*

Proof: The result follows from inspection of (4), combined with the preceding Steps. ∥ ▮

Proof of Lemma 1: The adviser's expected payoff from choosing p given the DM's belief p^e is given by:

$$\begin{aligned} U_A(p;p^e,B,\mu) &= p\Big[\mathbb{E}_{s\notin S(B,p^e)}\left[\mathbb{E}_\omega[u_1(\alpha_0(s),\omega)\mid s,\mu]\mid\mu\right]+\mathbb{E}_{s\in S(B,p^e)}\left[\mathbb{E}_\omega[u_1(a_\emptyset(B,p^e),\omega)|s,\mu]\mid\mu\right]\Big] \\ &\quad +(1-p)\mathbb{E}_\omega[u_1(a_\emptyset(B,p^e),\omega)|\mu]-c(p). \end{aligned}$$

The first term decomposes the adviser's payoff when he obtains the signal (with probability p): he reveals the signal if $s\notin S(B,p^e)$, which leads to the action $\alpha_0(s)$ by the DM; and he withholds the signal when $s\in S(B,p^e)$, which leads to the action $a_\emptyset(B,p^e)$ by the DM. The second term is the payoff when the adviser does not observe the signal (which arises with probability $1-p$), in which case the DM picks $a_\emptyset(B,p^e)$. The last term is the cost of information acquisition.

The conclusion of the Lemma follows from manipulating terms:

$$\begin{aligned} & U_A(p;p^e,B,\mu) \\ =\;& p\Big[\mathbb{E}_{s\notin S(B,p^e)}\left[\mathbb{E}_\omega[u_1(\alpha_0(s),\omega)\mid s,\mu]\mid\mu\right]+\mathbb{E}_{s\in S(B,p^e)}\left[\mathbb{E}_\omega[u_1(a_\emptyset(B;p^e),\omega)|s,\mu]\mid\mu\right]\Big] \\ & \quad +(1-p)\mathbb{E}_\omega[u_1(a_\emptyset(B;p^e),\omega)|\mu]-c(p) \\ =\;& p\left(\mathbb{E}_{s\notin S(B,p^e)}\left[\mathbb{E}_\omega[u_1(\alpha_0(s),\omega)\mid s,\mu]\mid\mu\right]+\mathbb{E}_{s\in S(B,p^e)}\left[\mathbb{E}_\omega[u_1(a_\emptyset(B;p^e),\omega)|s,\mu]\mid\mu\right]\right) \\ & \quad +(1-p)\mathbb{E}_s[\mathbb{E}_\omega[u_1(a_\emptyset(B;p^e),\omega)|s,\mu]]-c(p) \\ =\;& p\left(\mathbb{E}_{s\notin S(B,p^e)}\left[\mathbb{E}_\omega[u_1(\alpha_0(s),\omega)-u_1(a_\emptyset(B;p^e),\omega)\mid s,\mu]\mid\mu\right]\right) \\ & \quad +\mathbb{E}_s\left[\mathbb{E}_\omega[u_1(a_\emptyset(B;p^e),\omega)|s,\mu]\mid\mu\right]-c(p) \\ =\;& p\left(\int_{s\notin S(B,p^e)}\left[\left(a_\emptyset\left(B,p^e\right)-\rho s-B\right)^2-B^2\right]\gamma\left(s;\mu\right)ds\right) \\ & \quad -\int\left(a_\emptyset\left(B,p^e\right)-\left(\rho s+B\right)\right)^2\gamma\left(s;\mu\right)ds-\tilde{\sigma}^2-c(p), \end{aligned}$$

where the last expression is obtained from substituting (1). ▮

Proof of Lemma 2: By Proposition 1, there is full disclosure when $p^e=0$, hence we evaluate $\Delta(B,\mu,0)=\sigma_0^2+2\rho B>0$ from (7). For any (B,μ), $\Delta(B,\mu,\cdot):[0,\bar{p}]\to\mathbb{R}$ is a bounded mapping. Therefore, by the Inada conditions, we have $c'(0) = 0 < \Delta(B, \mu, 0)$ and $c'(p) > \Delta(B, \mu, 1)$ for large enough p. Since both sides of (8) are continuous in p, there exists $p \in (0, 1)$ that satisfies (8). It also follows that any equilibrium p must be interior, so it must satisfy (8). Finally, if p satisfies (8), we have

$$\frac{\partial U_A(\tilde{p}; p, B, \mu)}{\partial \tilde{p}} = \Delta(B, \mu, p) - c'(\tilde{p}) \gtreqless 0 \text{ if } \tilde{p} \lesseqgtr p,$$

due to the convexity of $c(\cdot)$, so p is an equilibrium effort choice. ▌

Proof of Proposition 3: This is a special case of Proposition 5. ▌

Proof of Proposition 4: Let $U(\mu)$ be the expected utility for the DM of appointing an adviser of prior μ. We can write

$$U(\mu) := p(\mu)W(\mu) + (1 - p(\mu))V(\mu),$$

where

$$\begin{aligned} W(\mu) &:= w(\mu, p(\mu)), \\ V(\mu) &:= v(\mu, p(\mu)), \end{aligned}$$

where

$$w(\mu, p) := -\tilde{\sigma}^2 - \int_{\underline{s}(B(\mu),p)}^{\overline{s}(B(\mu),p)} (a_\emptyset(B(\mu), p) - s\rho)^2 \gamma(s; 0) ds,$$

and

$$v(\mu, p) := -\tilde{\sigma}^2 - \int_{-\infty}^{\infty} (a_\emptyset(B(\mu), p) - s\rho)^2 \gamma(s; 0) ds.$$

We shall prove that $U'(0) = 0$ but the right second derivative of $U(\mu)$ evaluated at $\mu = 0$, denoted $U''(0^+)$, is strictly positive.[31] This will mean that the DM prefers an adviser with some $\mu > 0$ to an adviser with $\mu = 0$.

STEP 1. $p''(0) > p'(0) = 0$.

PROOF: Rewrite the equilibrium condition (8) for the adviser's effort choice as:

$$\mathcal{A}(\mu, p(\mu)) - c'(p(\mu)) = 0, \tag{14}$$

[31] A symmetric argument will prove that the left second derivative of U at $\mu = 0$ is strictly negative.

where $\mathcal{A}(\mu,p):=\Delta(B(\mu),\mu,p)$. Observe that

$$\begin{aligned}\mathcal{A}(\mu,p) &= \int_{s\notin S(B(\mu),p)} \rho(s-\overline{s}(B(\mu),p))\left(2(1-\rho)\mu+\rho(s-\overline{s}(B(\mu),p))\right)\gamma(s;\mu)ds \\ &= 2\mu\tilde{\sigma}^4\int_{s\notin S(B(\mu),p)}(s-\overline{s}(B(\mu),p))\gamma(s;\mu)ds \\ &\quad +\rho^2\int_{s\notin S(B(\mu),p)}(s-\overline{s}(B(\mu),p))^2\gamma(s;\mu)ds. \end{aligned} \tag{15}$$

A straightforward, but tedious, calculation yields

$$\mathcal{A}_{11}(0,p)>\mathcal{A}_1(0,p)=\mathcal{A}_2(0,p)=0, \tag{16}$$

where $\mathcal{A}_i$ denotes a partial derivative of $\mathcal{A}$ with respect to i-th variable, for $i=1,2$, and $\mathcal{A}_{ij}$ denotes a second partial derivative with respect to i and j-th variables for $i,j=1,2$.

There is a unique solution to (14) at $\mu=0$ because $\mathcal{A}(0,p)$ is a constant independent of p (since $\overline{s}(0,p)=\underline{s}(0,p)=0$) and $c'(p)$ is strictly increasing. Since $\mathcal{A}_2(0,p)=0<c''(p)$, the implicit function theorem implies that there is some neighborhood of $\mu=0$ where $p(\mu)$ is unique, continuously differentiable, and satisfies

$$p'(\mu)=\frac{\mathcal{A}_1(\mu,p(\mu))}{c''(p(\mu))-\mathcal{A}_2(\mu,p(\mu))}. \tag{17}$$

That $p''(0)>0$ is seen as follows:

$$\begin{aligned} p''(0) &= \lim_{\mu\to 0}\frac{p'(\mu)-p'(0)}{\mu} \\ &= \lim_{\mu\to 0}\frac{\mathcal{A}_1(\mu,p(\mu))}{\mu\left[c''(p(\mu))-\mathcal{A}_2(\mu,p(\mu))\right]} \\ &= \lim_{\mu\to 0}\frac{\mathcal{A}_{11}(\mu,p(\mu))+\mathcal{A}_{12}(\mu,p(\mu))p'(\mu)}{\mu\left[c'''(p(\mu))p'(\mu)-\mathcal{A}_{21}(\mu,p(\mu))-\mathcal{A}_{22}(\mu,p(\mu))p'(\mu)\right]+1\left[c''(p(\mu))-\mathcal{A}_2(\mu,p(\mu))\right]} \\ &= \frac{\mathcal{A}_{11}(0,p(0))}{c''(p(0))} \\ &> 0, \end{aligned}$$

where the second line uses (17) and $p'(0)=0$, the third line applies L'Hopital's rule, the fourth line again uses $p'(0)=0$, and the last line uses $\mathcal{A}_{11}(0,p(0))>0$.

Step 2. $W(0) - V(0) > 0$ and $W'(0) = V'(0) = W''(0) = V''(0) = 0$.

Proof: Recall $B(0) = 0$ and from Proposition 1 that $\overline{s}(0,p) = \underline{s}(0,p) = 0$. Hence, for any p,

$$w(0,p) - v(0,p) = \int_{-\infty}^{\infty} \left(a_{\emptyset}(0,p) - \frac{s\sigma_0^2}{\sigma^2} \right)^2 \gamma(s;0)ds > 0,$$

from which it follows that $W(0) - V(0) > 0$. For any p, direct computation yields

$$v_1(0,p) = v_2(0,p) = v_{11}(0,p) = w_1(0,p) = w_2(0,p) = w_{11}(0,p) = 0. \tag{18}$$

It follows from (18) and Step 1 that

$$\begin{aligned} W'(0) &= w_1(0,p(0)) + p'(0)\, w_2(0,p(0)) = 0, \\ V'(0) &= v_1(0,p(0)) + p'(0)\, v_2(0,p(0)) = 0, \\ W''(0) &= w_{11}(0,p(0)) + w_{12}(0,p(0))\, p'(0) \\ &\quad + p'(0)\left(w_{12}(0,p(0)) + w_{22}(0,p(0))\, p'(0)\right) + p''(0)\, w_2(0,p(0)) \\ &= 0, \end{aligned}$$

and

$$\begin{aligned} V''(0) &= v_{11}(0,p(0)) + v_{12}(0,p(0))\, p'(0) \\ &\quad + p'(0)\left(v_{12}(0,p(0)) + v_{22}(0,p(0))\, p'(0)\right) + p''(0)\, v_2(0,p(0)) \\ &= 0. \parallel \end{aligned}$$

From Step 1 and Step 2, we obtain

$$U'(0^+) = p'(0)(W(0) - V(0)) + p(0)W'(0) + (1 - p(0))V'(0) = 0,$$

and

$$\begin{aligned} U''(0^+) &= p''(0)(W(0) - V(0)) + 2p'(0)(W'(0) - V'(0)) + p(0)W''(0) + (1 - p(0))V''(0) \\ &= p''(0)(W(0) - V(0)) > 0. \end{aligned}$$

Combined, $U'(0^+) = 0$ and $U''(0^+) > 0$ imply that there exists $\mu > 0$ such that $U(\mu) > U(0)$. ▌

Proof of Proposition 5: Consider any pair (B, μ) and (B', μ') satisfying the hypothesized condition. It is without loss to assume $B' \geq 0$ and $\mu' \geq 0$. Further, since $(B(b, \mu), \mu)$ and $(B(-b, -\mu), -\mu)$ are payoff equivalent and thus generate the same incentive for the adviser, it is without loss to assume $\mu \geq 0$. The condition then reduces to $(0, 0) \leq (|B|, \mu) < (B', \mu')$. We focus on the case in which $B \geq 0$. As we will argue later, the case of $B < 0$ can be treated by the same argument applied twice, one for a shift from (B, μ) to $(-B, \mu)$, and another for a shift from $(-B, \mu)$ to (B', μ').

Let $p(B, \mu)$ be the (largest) p supported in equilibrium given an adviser with (B, μ). Suppose now an adviser with (B', μ') is chosen, but the DM *believes* that the adviser will continue to choose $p = p(B, \mu)$. We prove below that, given such a belief, the adviser with (B', μ') will choose strictly higher $p' > p(B, \mu)$. It will then follow that, since the adviser's best response correspondence is upper-hemicontinuous in the DM's belief (by the Theorem of Maxima), there must exist $p'' > p(\mu)$ such that p'' is supported under (B', μ'), which would imply that $p(B', \mu') > p(B, \mu)$.

To prove the statement, suppose to the contrary that the adviser with (B', μ') will find it optimal to choose $p' \leq p(\mu)$ given DM's belief that the adviser will choose $p(B, \mu)$. The disclosure sub-game following the effort choice p' is characterized by the pair $(S(B', p(B, \mu)),$ $a_\emptyset(B', p(B, \mu))$. By the first-order condition, we must then have

$$\Delta(B', \mu', p(B, \mu)) = c'(p') \leq c'(p(B, \mu)) = \Delta(B, \mu, p(B, \mu)). \tag{19}$$

For notational simplicity, let $S(\tilde{B}) := S(\tilde{B}, p(B, \mu))$, $\overline{s}(\tilde{B}) := \overline{s}(\tilde{B}, p(B, \mu))$, and $\underline{s}(B) :=$ $\underline{s}(\tilde{B}, p(B, \mu)) = \underline{s}(\tilde{B}, p(B, \mu))$, and let $\overline{s} := \overline{s}(B)$ and $\overline{s}' := \overline{s}(B')$.

The proof follows several steps.

Step 1. *The following inequality holds.*

$$\begin{aligned} \Delta(B', \mu', p(B, \mu)) \quad &\geq \quad \Pi(B', \mu') \\ &:= \quad \int_{s \notin S(B)} \left[\left(a_\emptyset(B', p(B, \mu)) - \rho s - B'\right)^2 - B'^2 \right] \gamma(s; \mu') ds. \end{aligned} \tag{20}$$

Proof: By picking a nondisclosure interval S, given his type μ, effort p, and the DM's nondisclosure action $a_\emptyset$, the adviser's expected utility is

$$\pi\left(S; B, \mu, p, a_\emptyset\right) := p \int_{s \notin S} \left[\left(a_\emptyset - \rho s - B\right)^2 - B^2 \right] \gamma\left(s; \mu\right) ds - \left(a_\emptyset - \mu - b\right)^2 - \sigma_0^2.$$

Thus, since the adviser chooses $S(B', p(B,\mu))$ rather than $S(B, p(B,\mu))$, it must be that

$$\pi\left(S(B', p(B,\mu)); B', \mu', p, a_{\emptyset}(B', p(B,\mu))\right) \geq \pi\left(S(B, p(B,\mu)); B', \mu', p, a_{\emptyset}(B', p(B,\mu))\right),$$

which implies the desired inequality by the definition of π. $\|$

Step 2. $\Pi(B', \mu') > \Pi(B, \mu)$.

Proof: By substituting for $a_{\emptyset}(B'; p(B,\mu)) = \rho\overline{s}(B'; p(B,\mu))$, we can write

$$\begin{aligned} \Pi(B', \mu') &= (2B' - 2\rho\overline{s}') \int_{s \notin S(B)} s\gamma(s;\mu')ds + \rho^2 \int_{s \notin S(B)} s^2\gamma(s;\mu')ds \\ &\quad + \left(\rho^2\overline{s}'^2 - 2B'\overline{s}'\right) \int_{s \notin S(B)} \gamma(s;\mu')ds. \end{aligned}$$

We then obtain the desired inequality:

$$\begin{aligned} & \Pi(B', \mu') \\ = \; & \Pr\{s \notin S(B) \mid \mu'\} \left\{2(B' - \rho\overline{s}')\mathbb{E}[s \mid s \notin S(B), \mu'] + \left(\rho^2\overline{s}'^2 - 2B'\overline{s}'\right)\right\} \\ & \quad + \rho^2 \int_{s \notin S(B)} s^2\gamma(s;\mu')ds \\ > \; & \Pr\{s \notin S(B) \mid \mu\} \left\{2(B - \rho\overline{s})\mathbb{E}[s \mid s \notin S(B), \mu'] + \left(\rho^2\overline{s}^2 - 2B\overline{s}\right)\right\} \\ & \quad + \rho^2 \int_{s \notin S(B)} s^2\gamma(s;\mu')ds \\ \geq \; & \Pr\{s \notin S(B) \mid \mu\} \left\{2(B - \rho\overline{s})\mathbb{E}[s \mid s \notin S(B), \mu'] + \left(\rho^2\overline{s}^2 - 2B\overline{s}\right)\right\} \\ & \quad + \rho^2 \int_{s \notin S(B)} s^2\gamma(s;\mu)ds \\ = \; & \Pi(B, \mu). \end{aligned}$$

The first inequality follows from the fact that $(B', \mu') > (B, \mu)$ and that $B \geq 0$. That $B' \geq B \geq 0$ implies $0 \geq \overline{s} = \overline{s}(B) \geq \overline{s}(B') = \overline{s}'$ (Step 3 in the proof of Proposition 1), which in turn implies that $B' - \rho\overline{s}' \geq B - \rho\overline{s} \geq 0$. Next, $\mathbb{E}[s \mid s \notin S(B), \mu'] \geq \mathbb{E}[s \mid s \notin S(B), \mu]$, since the Normal density $\gamma(\cdot;\mu')$ dominates in likelihood ratio the Normal density $\gamma(\cdot;\mu)$. We also have $\mathbb{E}[s \mid s \notin S(B), \mu'] \geq 0$, since $\mu' \geq 0$ (which follows from the fact that $|\mu'| \geq |\mu|$ and that $\mu' \geq \mu$) and since $S(B) \subset \mathbb{R}_{-}$. Next, $\mu' \geq \mu$ implies that $\Pr\{s \notin S(B) \mid \mu'\} \geq \Pr\{s \notin S(B) \mid \mu\}$. Combining all these facts imply the first inequality in weak form. The inequality is strict, however, since $B' > B$ or $\mu' > \mu$, which means one of the inequalities established above must be strict.

The second inequality is established as follows:

$$\begin{aligned}
\int_{s \notin S(B)} s^2 \gamma(s;\mu') ds &= \mathbb{E}[s^2|\mu'] - \Pr\{s \in S(B) \mid \mu'\} \, \mathbb{E}[s^2|s \in S(B), \mu'] \\
&= \mu'^2 - \sigma_1^2 - \Pr\{s \in S(B) \mid \mu'\} \, \mathbb{E}[s^2|s \in S(B), \mu'] \\
&\geq \mu^2 - \sigma_1^2 - \Pr\{s \in S(B) \mid \mu\} \, \mathbb{E}[s^2|s \in S(B), \mu] \\
&= \int_{s \notin S(B)} s^2 \gamma(s;\mu) ds,
\end{aligned}$$

where the inequality follows since $|\mu'| \geq |\mu|$, since $\mathbb{E}[s^2|\mu'] \leq \mathbb{E}[s^2|\mu]$ (which follows from the fact that $\gamma(\cdot;\mu')$ likelihood-ratio dominates $\gamma(\cdot;\mu)$ and that s^2 is decreasing in s for $s \in S(B) \subset \mathbb{R}_-$), and since $\Pr\{s \in S(B) \mid \mu'\} \leq \Pr\{s \in S(B) \mid \mu\}$. The string of inequalities thus proves the claim. $\|$

Combining Step 1 and Step 2, we have

$$\Delta(B', \mu', p(B,\mu)) > \Pi(B,\mu). \tag{21}$$

By definition, it also follows that

$$\Pi(B,\mu) = \Delta(B, \mu, p(B,\mu)). \tag{22}$$

Combining (21) and (22) yields

$$\Delta(B', \mu', p(B,\mu)) > \Delta(B, \mu, p(B,\mu)),$$

which contradicts (19). We have thus proven the statement of the proposition.

The case of $B < 0$ can be treated by applying the same sequence of arguments twice, one for a shift from (B,μ) to $(-B,\mu)$, and then another for a shift from $(-B,\mu)$ to (B',μ'). The second step satisfies the hypothesized condition, so the same argument works. The first step poses a slightly novel situation with Step 2. Yet, the same inequality works with $(B',\mu') := (-B,\mu)$. ▌

Proof of Proposition 6: The first statement follows from the discussion preceding the

Proposition. The second statement is a corollary of Proposition 4 and the observation that delegation and communication are equivalent when the adviser is like-minded. ▌

Proof of Proposition 7: First consider $\rho_A \geq \frac{\rho_{DM}}{2}$. It is straightforward to verify that if $\rho_A \geq \frac{\rho_{DM}}{2}$, there is a full disclosure equilibrium in the disclosure sub-game, independent of effort, p. Given full disclosure, the gain in utility for the adviser from observing a signal s over not observing it is

$$\int_{-\infty}^{\infty} (\rho_{DM} s - \omega)^2 \gamma(\omega|s, \rho_A) d\omega + \int_{-\infty}^{\infty} (-\omega)^2 \gamma(\omega|s, \rho_A),$$

where $\gamma(\cdot|s, \rho_A)$ denotes the posterior distribution over the state for the adviser with type ρ_A given signal s. The above can be simplified via algebra to

$$\rho_{DM} s^2 (2\rho_A - \rho_{DM}),$$

whose derivative with respect to ρ_A is $2\rho_P s^2$, which is strictly positive at all $s \neq 0$. Therefore, the marginal benefit of acquiring a signal is strictly higher for a more confident adviser, and consequently he exerts more effort. Since we have full disclosure in the disclose sub-game, it follows that the DM strictly prefers an adviser with $\rho_A = 1$ among all $\rho_A \geq \frac{\rho_{DM}}{2}$.

For $\rho_A < \frac{1}{2}\rho_{DM}$, it suffices to prove that there is a unique equilibrium in the disclosure sub-game, independent of effort, in which the adviser never discloses a signal. Consider any nondisclosure action $a_\emptyset \geq 0$ (the argument for $a_\emptyset < 0$ is symmetric to $a_\emptyset > 0$). It is straightforward to verify that the adviser's best response in terms of nondisclosure region is

$$S(a_\emptyset) = \left(-\infty, -\frac{a_\emptyset}{\rho_{DM} - 2\rho_A}\right] \cup \left[\frac{a_\emptyset}{\rho_{DM}}, \infty\right).$$

Note that this is no disclosure if (and only if) $a_\emptyset = 0$. It follows that $a_\emptyset = 0$ is an equilibrium independent of effort, p, because if the adviser is never disclosing, the DM will follow his prior upon nondisclosure. To see that there is no equilibrium with $a_\emptyset > 0$, observe that for any $a_\emptyset \geq 0$, the set $S(a_\emptyset)$ has expectation (with respect to the DM's prior density on signals) no greater than 0, and hence the DM's best response to $S(a_\emptyset)$, $a_N(p, S(a_\emptyset))$, is no larger than 0 for any p. Consequently, there does not exist $a_\emptyset > 0$ such that $a_\emptyset = a_N(p, S(a_\emptyset))$. ▌

References

Acemoglu, D., V. Chernozhukov, and M. Yildiz (2007): "Learning and Disagreement in an Uncertain World," mimeo, Massachusetts Institute of Technology. 3

Admati, A. R., and P. Pfleiderer (2004): "Broadcasting Opinions with an Overconfident Sender," *International Economic Review*, 45(2), 467–498. 24

Aghion, P., and J. Tirole (1997): "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105(1), 1–29. 6, 23

Aumann, R. J. (1976): "Agreeing to Disagee," *Annals of Statistics*, 4, 1236–1239. 3

Austen-Smith, D. (1994): "Strategic Transmission of Costly Information," *Econometrica*, 62(4), 955–63. 29

Banerjee, A., and R. Somanathan (2001): "A Simple Model Of Voice," *The Quarterly Journal of Economics*, 116(1), 189–227. 3, 7

Bergemann, D., and J. Valimaki (2002): "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 70(3), 1007–1033. 26

Blackwell, D. (1951): "Comparison of Experiments," *Proceedings of the Second Berkeley Symposium on Mathematical Statistics*, pp. 93–102. 7

Cai, H. (2004): "Costly Participation and Heterogeneous Preferences in Informational Committees," mimeo, UCLA.

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

Degroot, M. H. (1970): *Optimal Statistical Decisions.* McGraw-Hill, New York. 9

Dessein, W. (2002): "Authority and Communication in Organizations," *Review of Economic Studies*, 69(4), 811–838. 23

Dewatripont, M., and J. Tirole (1999): "Advocates," *Journal of Political Economy*, 107(1), 1–39. 26

Dixit, A. K., and J. W. Weibull (2007): "Political Polarization," mimeo, Princeton University and Stockholm School of Economics. 4

Dur, R., and O. H. Swank (2005): "Producing and Manipulating Information," *Economic Journal*, 115(500), 185–199. 25

Eliaz, K., and R. Spiegler (2006): "Contracting with Diversely Naive Agents," *Review of Economic Studies*, 71(3), 689–714. 7

Fudenberg, D., and J. Tirole (1991): *Game Theory*. MIT Press, Cambridge, MA. 8

Galilei, G. (1953): *Dialogue Concerning the Two Chief World Systems, Ptolemaic & Copernican*. Berkeley, University of California Press. 2

Gerardi, D., and L. Yariv (2007a): "Costly Expertise," mimeo, Yale University and Caltech. 25, 26

——— (2007b): "Information Acquisition in Committees," forthcoming, *Games and Economic Behavior*. 26

Gershkov, A., and B. Szentes (2004): "Optimal Voting Scheme under Costly Information Acquisition," mimeo, University of Chicago. 26

Green, J. R., and N. L. Stokey (2007): "A Two-person Game of Information Transmission," *Journal of Economic Theory*, 127(1), 90–104. 7

Grossman, S. J. (1981): "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law & Economics*, 24(3), 461–483. 6, 24

Harrington, Joseph E, J. (1993): "Economic Policy, Economic Performance, and Elections," *American Economic Review*, 83(1), 27–42. 7

Hori, K. (2006): "The Bright Side of Private Benefits," mimeo, Hitotsubashi University. 25

Janis, I. L. (1972): *Victims of Groupthink: A Psychological Study of Foreign-policy Decisions and Fiascoes*. Houghton Mifflin. 2

Kawamura, K. (2007): "Confidence and Competence in Expertise," mimeo, University of Oxford (Nuffield College). 24

Krishna, V., and J. Morgan (2001): "A Model of Expertise," *Quarterly Journal of Economics*, 116(2), 747–775. 25

Li, H. (2001): "A Theory of Conservatism," *Journal of Political Economy*, 109(3), 617–636. 26

Li, H., and W. Suen (2004): "Delegating Decisions to Experts," *Journal of Political Economy*, 112(S1), S311–S335. 25

Matthews, S., and A. Postlewaite (1985): "Quality Testing and Disclosure," *RAND Journal of Economics*, 16(3), 328–340. 25

Milgrom, P., and J. Roberts (1986): "Relying on the Information of Interested Parties," *RAND Journal of Economics*, 17(1), 18–32. 8

Milgrom, P. R. (1981): "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12(2), 380–391. 6, 8, 24

Morris, S. (1995): "The Common Prior Assumption in Economic Theory," *Economics and Philosophy*, 11, 227–253. 7

Prendergast, C. (2007): "The Motivation and Bias of Bureaucrats," *American Economic Review*, 97(1), 180–196. 26

Shavell, S. (1994): "Acquisition and Disclosure of Information Prior to Sale," *RAND Journal of Economics*, 25(1), 20–36. 25

Shin, H. S. (1994): "The Burden of Proof in a Game of Persuasion," *Journal of Economic Theory*, 64(1), 253–264. 25

——— (1998): "Adversarial and Inquisitorial Procedures in Arbitration," *RAND Journal of Economics*, 29(2), 378–405. 8, 25

Spector, D. (2000): "Rational Debate and One-Dimensional Conflict," *Quarterly Journal of Economics*, pp. 181–200. 7

Suen, W. (2004): "The Self-Perpetuation of Biased Beliefs," *Economic Journal*, 114(495), 377–396. 25

Szalay, D. (2005): "The Economics of Clear Advice and Extreme Options," *Review of Economic Studies*, 72, 1173–1198. 23, 26

Tversky, A., and D. Kahneman (1974): "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185, 1124–1131. 3

Van den Steen, E. (2004): "The Costs and Benefits of Homogeneity, with an Application to Culture Clash," mimeo, MIT Sloan. 26

Van den Steen, E. (2005): "Organizational Beliefs and Managerial Vision," *Journal of Law, Economics and Organization*, 21(1), 256–283. 7

Yildiz, M. (2003): "Bargaining without a Common Prior—An Immediate Agreement Theorem," *Econometrica*, 71(3), 793–811. 7

——— (2004): "Waiting to Persuade," *Quarterly Journal of Economics*, 119(1), 223–248. 15