STRATEGIC ARGUMENTATION[1]

Wioletta Dziuda
Kellogg School of Management
Managerial Economics and Decision Sciences
wdziuda@kellogg.northwestern.edu

January 26, 2009

[1]I thank Dilip Abreu, Paolo Ghirardato, Faruk Gul, Johannes Hörner, Navin Kartik, Massimo Marinacci, Eric Maskin, Giovanni Mastrobuoni, Stephen Morris, Marco Ottaviani, Wolfgang Pesendorfer, Yuval Salant, Hyun Shin, Jonathan Vogel, and especially my adviser, Markus Brunnermeier, for helpful comments. Seminar participants at Bocconi University, Collegio Carlo Alberto, Harvard University, IESE Business School, INSEAD, Kellogg School of Management, London School of Economics, Princeton University, University of Maryland, University of Minnesota and Washington University in St. Louis provided further useful comments.

Abstract

I analyze a disclosure game between an uninformed decision maker and an informed but possibly biased expert. The relevant information is contained in a set of arguments. The expert can disclose each argument credibly, but he cannot prove whether he has disclosed everything.

In all equilibria some, but not all, information is revealed. The biased expert exaggerates his reports in favor of his preference, yet he does not suppress all of the unfavorable information. The decision maker takes balanced reports at face value, but is skeptical about the unbalanced ones.

The model captures the use of two-sided messages and can explain occurrence of confirmatory bias.

Keywords Strategic communication, persuasion, argumentation, expert, disclosure games

1 Introduction

Consider a first-time camera buyer who is uninformed about the complexity of the product; she does not know which and how many technical specifications are important for taking quality pictures. The salesperson knows all the relevant features and can credibly disclose each of them to the buyer. Disclosed information is credible because of different reasons: the salesperson may be able to prove the size of the display and zoom by demonstrating them, the consumer may be able to test the features by using the camera, or liability laws may make it unprofitable for the salesperson to lie about any feature. Drawing upon everyday experience, one might expect two things to happen in this interaction. First, the salesperson may not fully inform the buyer. This observation is not in line with the unravelling result, which states that when the expert cannot lie, he or she reveals all information even if the parties have conflicting interests.[1] Second, the salesperson may not completely suppress the unfavorable information about the camera; in addition to revealing its favorable characteristics, the salesperson may also mention some unfavorable ones—for example, he or she may mention its short battery life.

There are many situations in which the expert does not disclose all information, but presents arguments which interpreted at face value would go against what the expert is arguing for: a financial adviser informing a client about investment options, a doctor advising a patient on the selection of a treatment, or an author of an article trying to influence the reader. The goal of this paper is to study whether these observations can be consistent with a game theoretic model.

The model has the following structure. A decision maker (she) consults an expert (he) to help her choose between two alternatives. The decision maker prefers the alternative which is favored by a sufficient fraction of arguments, but she is uninformed about the number of arguments of each type. The expert receives a sequence of arguments, each of which favors one alternative, and he can disclose credibly any argument. The expert may be either an honest type who reveals all of the arguments, or he may try to convince the decision maker to choose a particular alternative.

Consistent with the camera buyer example, full disclosure of information is not an equi-

[1]Since Grossman (1981), Milgrom (1981), Milgrom and Roberts (1986), and Matthews and Postlewaite (1985) first laid out the unravelling argument, the research has focused on identifying situations in which this argument may fail. Viscusi (1978) and Jovanovic (1982) show that the expert reveals only favorable states when disclosure is costly. In Fishman and Hagerty (2003), information transmission is hindered by the presence of decision makers who can verify the event of disclosure, but do not understand the disclosed information. Dye (1985), Shin (1994a), Shin (1994b), and Shin (2003) show that unravelling may fail if there is uncertainty about how well informed the expert is. In a dynamic setup, Grubb (2007) shows that unravelling may be hindered further if senders want to build a reputation for being uninformed.

librium; if it were, the decision maker would take all reports at face value, but then a biased expert would have an incentive to reveal only arguments favorable to him. Unravelling fails because the decision maker does not know how many arguments she should expect from the expert, and the expert is unable to prove how many arguments exist.

The main result of the paper is that in any equilibrium the biased expert indeed reveals some unfavorable arguments. The reason for this is as follows. If in equilibrium the biased expert always concealed all unfavorable arguments, the decision maker would discount such messages accordingly, but take all other messages at face value. In such a case, revealing at least one unfavorable argument could allow the biased expert to fool the decision maker, which cannot happen in equilibrium. This finding can be interpreted as a use of two-sided messages – messages containing arguments both for and against a given alternative –which is consistent with everyday experience. The abovementioned camera salesperson may mention some negative features of the product, a car dealer may mention a car's long acceleration time, or an author may mention findings that disagree with his agenda. Some commercials use two-sided messages, for example, an advertisement attempting to persuade consumers of superiority of dBase IV software disclosed that it was more costly and worse at handling errors than competing products.[2]

The second finding of the paper is that despite the equilibrium use of unfavorable arguments, unless the expert presents a very balanced set of arguments, the decision maker bases her decision only on the arguments favoring the alternative that the expert is arguing for. She disregards any unfavorable arguments that the expert reveals! The reason for this is as follows. Given that a biased expert can conceal unfavorable arguments even when his preferred alternative is not very attractive, the decision maker must make the same choice independent of how many unfavorable arguments the expert reveals. If she did not, the biased expert would mislead her by revealing enough unfavorable arguments that would make the decision maker choose the expert's preferred alternative.

It is worth emphasizing that the use of two-sided messages in this model does not arise from dynamic considerations, such as reputation or career concerns. The rationale for the use of two-sided messages can be better understood if we think of an economy with infinitely many experts and decision makers who meet for one-shot interactions. Clearly, the decision makers must be skeptical whenever they hear an unbalanced set of arguments, otherwise they would be susceptible to manipulation. Their beliefs, however, must be sustained by the behavior of the experts. If the experts stopped using two-sided messages, the decision

[2]See Pechmann (1992). Another example is the case of Continental Airlines acknowledging a variety of past problems such as canceled flights and lost luggage when trying to persuade the clients about its new commitment to quality (Crowley and Hoyer, 1994).

makers would update their beliefs accordingly, and the experts would have a strict incentive to use them. [3]

A surprising result of the current model is that it can generate confirmatory bias. Confirmatory bias, which has been regarded as a non-Bayesian phenomenon, is an observation that when confronted with mixed, ambiguous evidence, people tend to become even more extreme in their beliefs. In my model, the decision maker presented with unbalanced evidence bases her decision only on the arguments that favor the alternative that the expert is arguing for. She believes the quality of this alternative is higher, the more favorable arguments she hears. When presented with additional, mixed evidence, she updates her belief based only on the favorable arguments, and therefore she becomes more optimistic about this alternative.

The model sheds some light on a number of interesting issues. For example, it allows to address the question whether mandating the disclosure of the conflict of interest benefits the decision maker. The model says that if the problem which the expert advises on is simple, that is, is characterized by few arguments only, the expert would like to hide his potential bias. In such situations, mandating disclosure could increase the decision maker's utility. If the issue is complex, however, such disclosure is inconsequential. This might explain why in practice experts sometimes voluntary disclose their associations that create the conflict of interest.

Complexity of the problem affects also how much information is disclosed. When the problem is simple, the decision maker might choose a suboptimal alternative even if all types of the persuader would like to convince her otherwise. When the complexity of the problem increases, the decision maker is more susceptible to manipulation. However, the honest expert might be more able to convey information when the problem is complex.

This paper complements extensive literature on communication and disclosure games.[4] The papers closest to mine, Shin (1994a) and Wolinsky (2003), study games with similar relation between the payoff-relevant state and the set of available messages. Shin (1994a)

[3] Psychological research has shown that two-sided messages are more persuasive and increase the perceived truthfulness of the expert (see, for example, Smith and Hunt,1978; Anderson and Golden,1984). In the model, two-sided messages are not more persuasive, but neither do they harm the expert, and in equilibrium the persuader must use them to avoid revealing his type. However, adding a small number of naive decision makers who take the expert's messages at face value makes the rational decision maker more likely to be persuaded by the expert using a two-sided message. In such a game, the persuader has an incentive to bias his reports to influence the naive decision makers, which is why a rational decision maker is more skeptical about reports with few unfavorable arguments. This result is available from the author upon request.

[4] Among many others, Crawford and Sobel (1982), Krishna and Morgan (2001), Battaglini (2002), Chen (2007) analyze communication situations in which talk is cheap. Kartik (2008) and Kartik, Ottaviani and Squintani (2007) analyze communication games in which misrepresenting the state is costly. Verrecchia (2001) is an extensive survey of the literature on disclosure games. From the modeling perspective, Glazer and Rubinstein (2001) and (2004) also analyze a model in which information is a collection of arguments, but they focus on communication when there is boundary on the number of arguments that can be revealed.

analyzes a game in which the expert may have imprecise information about the state of nature and can reveal what he knows credibly, but he cannot prove how imprecise his information is. Our models differ in details, but the rationale for the lack of information revelation is somewhat similar: the expert is unable to prove whether he has disclosed everything. The most important difference between the current paper and Shin (1994a) is that Shin (1994a) does not allow for the uncertainty about expert's preferences and focuses on the class of equilibria in which the expert reveals only favorable information. As the current paper shows, allowing for uncertainty about the preferences of the expert, makes the complete suppression of unfavorable information not an equilibrium.[5] Wolinsky (2003) introduces uncertainty about the preference of the expert in a game with fully informed experts. In Wolinsky (2003) and the current paper, the uncertainty about the expert's preferences results in the use of reports that would not be most favorable if taken at face value. In contrast to this paper, however, in Wolinsky (2003) the strategy of the expert cannot be interpreted as the use of two–sided messages.

The nature of disclosure suggests that the message space available to the expert is discrete, and this is what Shin (1994a), Shin (1994b), and Wolinsky (2003) assume. One technical contribution of this paper is to model the number of arguments as a continuous variable, which makes the model more tractable.

The paper is organized as follows. Section 2 describes the game. Section 3 analyzes a version of the model in which the expert is either honest or biased toward *Right.* Section 4 extends the analysis to the case in which the expert can be any of three types. Section 5 discusses how confirmatory bias may arise in the model. Section 6 analyzes the impact of uncertainty that the decision maker faces on the equilibrium outcome. Section 7 provides a summary and conclusions.

2 The Model

The environment

There are two alternatives: *Right* and *Left.* A state of nature is a tuple $(L, R) \in R^2_+$. L represents the number of arguments in favor of *Left,* and R represents the number of arguments in favor of *Right.* Let $f(L, R)$ be the prior probability density function over the state space, and $F(L, R)$ be the corresponding distribution function. The distribution over the state space is common knowledge.[6]

[5] For more on this difference, see Section 6.1.

[6] Note that arguments are continuous variables. In my motivating examples, the number of arguments is a discrete variable, but modeling L and R as continuous variables makes the model significantly more tractable and allows it to obtain general results without putting much structure on $f(L, R)$. Section 6.3

There are two players: an expert and a decision maker.

The expert

The expert observes the state of nature (L, R) and sends a *report* to the decision maker. A report is a tuple (λ, ρ), where λ is the number of arguments in favor of *Left* and ρ is the number of arguments in favor of *Right* that the expert reveals. A report $(\lambda, \rho) \in R_+^2$ is *feasible* in state (L, R) if $\lambda \leq L$ and $\rho \leq R$, and in each state the expert can send any report from the feasible set at no cost. That is, the expert can truthfully disclose any subset of the existing arguments, but cannot credibly convey that he has disclosed all of them.

There are three types of experts: an honest expert, H; a persuader toward *Right,* P_r; and a persuader toward *Left*, P_l. An honest expert is *non-strategic* and reveals all of the arguments. A persuader toward alternative A wants the decision maker to choose $A,$ independent of the state of nature; that is, he maximizes $\Pr\{A \text{ is chosen}\}$. The probability that the expert is of type $i \in \{P_l, P_r, H\}$ is π_i. A strategy of a type i persuader, $m_i((\lambda, \rho) \,|\, (L, R))$, specifies for each (L, R) the probability distribution over the set of feasible reports. The expert is said to report *fully* if he reveals all of the arguments.

The decision maker

Define the quality of a given alternative as the fraction of arguments in its favor, $q_R = \frac{R}{R+L}$ and $q_L = \frac{L}{R+L}$. The utility of the decision maker is:

$$U(Right|L, R) = q_R - \theta, \; U(Left|L, R) = q_L + \theta - 1 \tag{1}$$

where $\theta \in [0, 1]$ is a preference parameter. In short, the decision maker chooses *Right* if and only if $E[q_R|\lambda, \rho] \geq \theta$.

The parameter θ describes an ex ante preference of the decision maker. For example, a consumer may have some intrinsic preference for Canon cameras over Panasonic cameras, a shareholder may prefer stocks of environmentally friendly companies, or a voter may prefer a Republican candidate because of family tradition, all other things being equal. Nature chooses θ according to a continuous probability density function $g(\theta)$ with full support, with the corresponding distribution function $G(\theta)$. The decision maker observes her θ, but the expert does not.

The game

The game proceeds as follows. First, nature determines the type of expert $i \in \{P_l, P_r, H\}$, the type of the decision maker θ, and the set of arguments (L, R). The expert observes his type and the state of nature (L, R), and sends a report (λ, ρ) to the decision maker. The decision maker observes her type and the report, and chooses one of the alternatives.

discusses the limitations of this assumption.

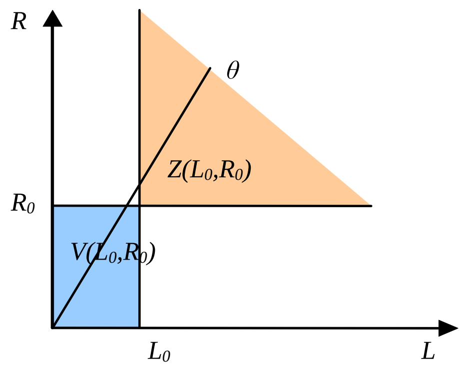

Figure 1: This figure represents the state space. $V(L_0, R_0)$ is the set of feasible reports for (L_0, R_0) and $Z(L_0, R_0)$ is the set of states of nature in which report (L_0, R_0) is feasible.

Assumptions on $f(L, R)$

Assume that $f(L, R)$ is continuous with full support on R_+^2. Let $f^L(L|R)$ denote the conditional density of L given R, and $f^R(R|L)$ the conditional density of R given L. Let $F^L(L|R)$ and $F^R(R|L)$ be the corresponding distribution functions. I impose the following regularity conditions:

$$\frac{dF^L(L|R)}{dR} \geq 0, \ \frac{dF^R(R|L)}{dL} \geq 0. \tag{2}$$

Intuitively, condition 2 says that the presence of an additional argument in favor of one alternative does not make the opposing arguments more likely. This rules out situations in which "good news is bad news," that is, in which a higher number of arguments favoring a given alternative makes this alternative less likely to be attractive.

Graphic representation

The triangle in Figure 1 represents the state space and the report space. Define $Z(\lambda, \rho) = \{(L, R) \in S : L \geq \lambda \text{ and } R \geq \rho\}$ and $V(L, R) = \{(\lambda, \rho) : \lambda \leq L \text{ and } \rho \leq R\}$. If the state of nature is (L_0, R_0), then the shaded region $V(L_0, \rho_0)$ is the set of feasible reports. The shaded region $Z(L_0, R_0)$ is the set of states of nature that allow the expert to send a report $(\lambda, \rho) = (L_0, R_0)$.

The line θ represents the states of nature that generate the same quality, $q_R = \theta$. The decision maker of type θ prefers to choose *Right* if the state of nature lies above this line but *Left* otherwise.

Equilibrium concept

I look for a perfect Bayesian equilibrium of this game. The decision maker in this game has a very limited role: she chooses *Right* if she believes that the quality of *Right* is above θ. Given this, from the perspective of the expert the probability that the decision maker chooses *Right* is a strictly increasing function of her belief. Therefore, one can transform this game into a game with one player only, in which the persuader toward A maximizes the belief of the decision maker about q_A, where the belief satisfies Bayes' rule.[7]

Let $\eta(\lambda,\rho)$ be the equilibrium belief of the decision maker about q_R if she observes report (λ,ρ). A perfect Bayesian equilibrium is characterized by m_i for $i \in \{P_l, P_r\}$ and $\eta(\lambda,\rho)$ such that

1. m_i satisfies
$$\int_0^R \int_0^L m_i\left((\lambda,\rho) \mid (L,R)\right) d\lambda d\rho = 1 \text{ for all } (L,R) \in R_+^2;$$
2. if (λ^*,ρ^*) is in the support of $m_{P_r}(\cdot \mid (L,R))$, then (λ^*,ρ^*) solves
$$\max_{(\lambda,\rho)\in V(L,R)} \eta(\lambda,\rho);$$
and if (λ^*,ρ^*) is in the support of $m_{P_l}(\cdot \mid (L,R))$, then (λ^*,ρ^*) solves
$$\min_{(\lambda,\rho)\in V(L,R)} \eta(\lambda,\rho);$$
3. $\eta(\lambda,\rho)$ satisfies Bayes' rule.

The above definition entails the immediate conclusion that the set of equilibria does not depend on the distribution of θ as long as it has full support.

3 One-sided Bias

In this section, I consider a situation in which the decision maker knows the direction of the potential bias of the expert. Let π be the probability that the expert is biased toward *Right,* i.e., is of type P_r; and $1-\pi$ be the probability that the expert is honest, i.e., is of type H.

The case with one type of persuader allows to present the main findings of the model in a simpler setting. It also describes better the situations in which the decision maker knows which alternative the expert may favor. A sales representative may advise the customer honestly about the quality of his product, but he is certainly not interested in increasing the

[7] The honest expert uses all possible reports; therefore, there are no off-equilibrium beliefs. However, given that there is a continuum of possible reports, in equilibrium some reports might be sent with probability 0, in which case $\eta(\lambda,\rho)$ is not well defined. However, $\eta(\lambda,\rho)$ must satisfy Bayes' rule on all sets of non-zero measure.

sales of competing products. A firm may report honestly, but if it does not, it is interested in maximizing its perceived value.

3.1 The properties of the equilibria

It is easy to see that there is no equilibrium with full information disclosure. If there were one, the expert's reports would be taken at face value, $\eta(\lambda, \rho) = \frac{\rho}{\lambda+\rho}$, and the persuader would prefer to conceal all unfavorable arguments, convincing the decision maker to choose *Right.* Proposition 1 states the less obvious result.

Proposition 1 *There is no equilibrium in which the decision maker's belief is independent of the expert's report (babbling equilibrium).*

Proof All proofs are in the Appendix. ■

The intuition for Proposition 1 is as follows. In a babbling equilibrium the decision maker must keep her prior belief $E\left[\frac{R}{R+L}\right]$ after all reports of the expert. Consider a decision maker who receives many arguments in favor of *Right* (ρ_0 large) and few in favor of *Left* (λ_0 small), such that $\frac{\rho_0}{\lambda_0+\rho_0} > E\left[\frac{R}{R+L}\right]$. There is some probability that these reports come from the honest expert, and in such a case the quality of *Right* is close to 1. Hence, if the decision maker belief is $E\left[\frac{R}{R+L}\right] < 1$, it must be that the persuader sends this report when the quality of *Right* is low, which is when there exist many more arguments in favor of *Left* than in the report. But given condition 2, such a state is very unlikely relative to state (λ_0, ρ_0). Hence for ρ_0 large enough the decision maker believes it is very likely that the report came from the honest expert; her belief therefore must be close to 1.

Although there is no babbling equilibrium, there are many partially revealing equilibria in this game, which is not very surprising given that the message space is larger than the payoff-relevant space. However, all equilibria share properties described in Proposition 2.

Proposition 2 *In each equilibrium, for all $\rho > 0$, there exists $\lambda_\rho > 0$ such that $\eta(\lambda, \rho)$ is constant for all $\lambda \in [0, \lambda_\rho)$, and this belief is weakly increasing in ρ. The persuader is indifferent among revealing any small number of unfavorable arguments and does not always suppress them completely.*

Proposition 2 contains two main findings of this paper. First, in all equilibria the persuader might not reveal all arguments but does not completely suppress the unfavorable information: his strategy includes revealing some arguments that appear to oppose his interest. Second, two-sided messages neither benefit nor harm the expert. If the revealed arguments are not balanced enough, the decision maker bases her decision solely on the number of arguments in favor of *Right:* $\eta(\lambda,\rho)$ is independent of λ.

The intuition for Proposition 2 is as follows. The decision maker should be skeptical about reports with very few unfavorable arguments, that is, she should form a lower belief than she would if she took them at face value. This is because the persuader can send such reports almost independent of the quality of *Right.* Moreover, if she forms different belief after hearing two different reports of this type, the persuader would only use the one that results in higher belief, and her skepticism would not be rational. Hence, she forms the same belief upon seeing any small number of unfavorable arguments. This means that revealing a small number of unfavorable arguments does not hurt the persuader, and in equilibrium he uses these arguments in order to keep the decision maker's beliefs justified.

Proposition 2 states that all equilibria of this model share many qualitative features. In fact, the main difference across the equilibria is the number of favorable arguments that the expert reveals. Contrary to what one might expect, the expert might not reveal all of the favorable arguments. This can happen if the decision maker expects the expert to conceal some of the favorable arguments in states that are very good for the expert. Although interesting, such behavior does not seem very realistic. In these equilibria, the decision maker's belief is discontinuous in the number of observed arguments, and it is difficult to understand why such belief would arise. On the other hand, in all equilibria in which the expert reveals all favorable arguments, the belief function $\eta(\lambda,\rho)$ is a continuous function of the expert's report. Moreover, all such equilibria are outcome equivalent. For these reasons as well as for expositional simplicity, in the reminder of this paper I focus only on the equilibria with the continuous belief function.[8]

Proposition 3 *There is a unique equilibrium belief function $\eta(\lambda,\rho)$ that is continuous in reports (λ,ρ). All equilibria characterized by this function are outcome-equivalent. In any such equilibrium, $\eta(\lambda,\rho)$ is strictly increasing in ρ, and for each ρ, there exists $\lambda_\rho > 0$, defined by*

$$\frac{\rho}{\rho+\lambda_\rho} = \Pr\left(H|\lambda \le \lambda_\rho, \rho\right) E\left[q_R|L \le \lambda_\rho, R=\rho\right] + \Pr\left(P_r|\lambda \le \lambda_\rho, \rho\right) E\left[q_R|R=\rho\right] \tag{3}$$

such that

i. P_r reveals all arguments in favor of Right*, $\rho = R$ for each R;*

[8] Section 6.3 shows that these are also the most informative equilibria. This could further justify selection of these equilibria.

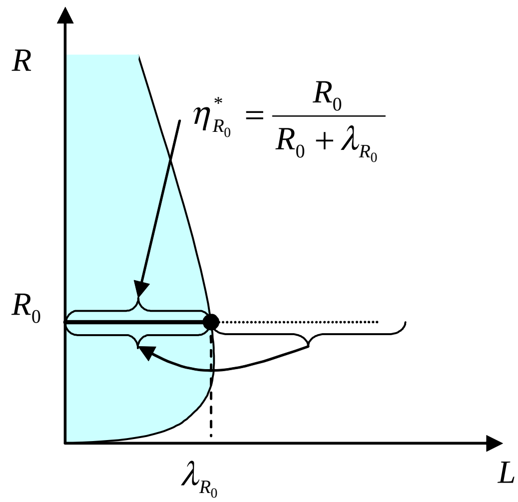

Figure 2: The details of the continuous equilibria. The shaded region represents the ambiguity area.

ii. *P_r reveals a subset of arguments in favor of* Left*,* $\lambda \leq \min\{L, \lambda_\rho\}$ *, using a strategy that results in* $\eta(\lambda, \rho) = \frac{\rho}{\rho + \lambda_\rho} \equiv \eta^*_\rho$ *for all* $\lambda \leq \lambda_\rho$*.*

Figure 2 represents all of these equilibria for a fixed π and $f(L, R)$.[9] The white area, which I call the *revealing area*, is the set of reports that are used in equilibrium only by the honest expert. The shaded region, which I call the *ambiguity area,* is the set of reports used in equilibrium also by the persuader. Hence, the ambiguity area includes all reports that do not allow the decision maker to identify the type of expert. The boundary of the ambiguity area is determined by $\lambda_{\rho=R}$ defined by equation (3).

In any continuous equilibrium, the persuader reveals all of the arguments that favor *Right* and some arguments that favor *Left.* The highest number of arguments in favor of *Left* that the persuader reveals for any R is λ_R. After observing a report from the revealing area, the decision maker knows that the expert has reported fully, and she takes this report at face value, that is, $\eta(\lambda, \rho) = \frac{\rho}{\rho + \lambda}$. After observing a report from the ambiguity area, she forms her belief based only on $\rho : \eta(\lambda, \rho) = \frac{\rho}{\rho + \lambda_\rho}$. Given this belief function, the persuader is indifferent among sending any report (λ, ρ) such that $\rho = R$ and $\lambda \leq \lambda_R$. However, in equilibrium he must use a strategy that supports the decision maker's beliefs.

The persuader can use many strategies that support the continuous belief function, but they all must generate a belief that is constant in λ for $(\lambda, \rho) : \lambda \leq \lambda_\rho$. This belief is

[9] The comparative statics is in Section 6.

characterized by

$$\eta(\lambda,\rho) = \Pr(H|\lambda,\rho)\frac{\rho}{\rho+\lambda} + (1-\Pr(H|\lambda,\rho))\, E\left[\frac{\rho}{\rho+L}|\lambda,\rho,P_r\right]. \tag{4}$$

If the expert is honest, then a smaller number of arguments in favor of *Left* implies a higher quality of *Right.* The strategy of the persuader must offset this effect. Differentiating both sides of equation 4 with respect to λ, we get

$$\begin{aligned} & \frac{d\Pr(H|\lambda,\rho)}{d\lambda}\left(\frac{\rho}{\rho+\lambda} - E\left[\frac{\rho}{\rho+L}|\lambda,\rho,P_r\right]\right) \\ = \quad & \Pr(H|\lambda,\rho)\frac{\rho}{(\rho+\lambda)^2} - (1-\Pr(H|\lambda,\rho))\frac{dE\left[\frac{\rho}{\rho+L}|\lambda,\rho,P_r\right]}{d\lambda}. \end{aligned}$$

The expression in brackets on the left-hand side is positive because the persuader sends λ only if $L \geq \lambda$. This implies that in equilibrium either $\frac{d\Pr(H|\lambda,\rho)}{d\lambda} > 0$ or $\frac{dE\left[\frac{\rho}{\rho+L}|\lambda,\rho,P_r\right]}{d\lambda} > 0$. This means that as the number of the unfavorable arguments that the expert reveals increases, either the posterior belief that the expert is honest increases, or the expected quality of *Right* conditional on the expert being a persuader increases – or both. The first effect captures the intuition that two-sided arguments might be more credible (although not more persuasive). The second effect is less intuitive; it says that the less favorable the state is, the more likely the persuader is to send more extreme reports . Although there is no guarantee that revealing unfavorable information increases the credibility of the expert in any particular equilibrium, there always exist equilibria in which this is true. In one such equilibrium the persuader reports fully if $L \leq \lambda_\rho$; otherwise, he randomizes over $\lambda \in [0,\lambda_\rho]$ using some probability density function $s(\lambda)$. For $s(\lambda)$ to be an equilibrium, it must be decreasing in λ, which means that the persuader randomizes over how many arguments to reveal, but is more likely to reveal fewer of them. Under this strategy revealing an additional argument in favor of *Left* increases the credibility of the expert, $\Pr(H|\lambda,\rho)$, but decreases the estimate of the quality of *Right* conditional on the expert being honest.

Figure 3 shows the behavior of the decision maker of type θ in any continuous equilibrium. The first triangle represents her decision as a function of state if she happens to face the honest expert, and the second triangle represents her decision if she happens to face the persuader. The shaded areas represent the states in which the decision maker chooses *Left.* When the number of arguments that favor *Right* is low enough, the decision maker chooses *Left* even if she receives an extreme report.

Figure 3 shows that the decision maker can make suboptimal choices even in some situations in which the alternative preferred by all types of the expert coincides with the best

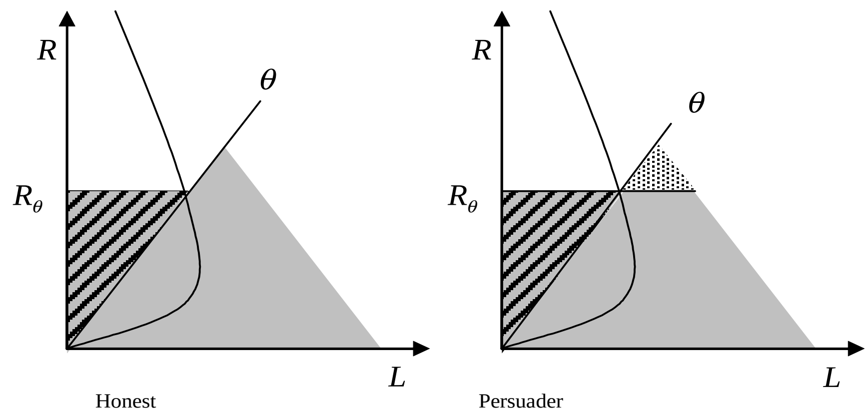

Figure 3: The behavior of the decision maker. The triangles represent the choice of the decision maker given the state of nature and given that the expert happens to be an honest type or a persuader, respectively.

alternative for the decision maker.[10] This happens when the sufficient majority of arguments support *Right,* but there are few arguments in total. The number of arguments in each state of the world can be interpreted as complexity of the problem; for example, products with more relevant characteristics such as cameras are more complex than products with few characteristics; issues that require looking at many aspects are more complex than issues where only a few facts matter. Under such notion of complexity, figure 3 shows that the choice of the decision maker is suboptimal from the point of view of all players when the problem turns out to be relatively simple and the biased expert's alternative is the better one (striped area). One can also see that both types of the expert prefer complex situations. When the complexity of the problem increases, the decision maker is more susceptible to manipulation (dotted area). However, the honest expert might be more able to convey information when the problem is complex.

4 Two-sided Bias

In this section, I consider the situation in which the expert can be biased toward either alternative. The expert can be P_r, P_l, or H with probabilities π_r, π_l, and π_H, respectively.

Sometimes the decision maker is uncertain not only about whether the expert is honest, but also about the potential bias of the persuader. A salesperson may give honest advice, or he may have an interest in selling one particular product, and the decision maker may not

[10] It would be natural to assume that the honest expert wants to maximize the utility of the decision maker.

know which product that is. Similarly, a scientist publishing a comparison of the performance of two drugs may be honest or biased, and the reader may not know which pharmaceutical company funded the research.

The previous section shows that when the expert is either honest or biased toward *Right*, the decision maker knows that all or majority of the arguments in favor of *Right* have been revealed and uses these arguments to form her beliefs. Unless the expert reveals himself to be honest by presenting many unfavorable arguments, she completely disregards the arguments that favor *Left.* When the expert can be biased in either direction, as it is the case in this section, the decision maker cannot use the same logic; therefore, we can expect that much less information will be revealed. This is, however, only partially true. Proposition 4 describes the unique continuous equilibrium outcome. The persuader toward *Right* and the persuader toward *Left* separate themselves if they happen to receive many arguments in favor of their preferred alternatives. In these states, the decision maker can use the same skeptical approach to infer information as in the one-sided case. If the persuaders receive very few favorable arguments, however, little information is revealed.

Proposition 4 *There is a unique equilibrium belief function $\eta(\lambda, \rho)$ that is continuous in reports (λ, ρ), and all equilibria characterized by this function are outcome-equivalent. In any such equilibrium, $\eta(\lambda, \rho)$ is weakly increasing in ρ and weakly decreasing in λ, and there exist unique thresholds $\bar{L}$ and $\bar{R}$ and unique functions λ_ρ and ρ_ρ such that:*

i. For all $\rho > 0$, $\lambda_\rho > 0$ and $\eta(\lambda, \rho)$ is constant for all $\lambda \in [0, \lambda_\rho]$.

ii. For all $\lambda > 0$, $\rho_\lambda > 0$ and $\eta(\lambda, \rho)$ is constant for all $\rho \in [0, \rho_\lambda]$.

iii. For each ρ, the persuader toward Right *reveals at most λ_ρ arguments in favor of* Left, *and for each λ, the persuader toward* Left *reveals at most ρ_λ arguments in favor of* Right.

iv. For $R \geq \bar{R}$ the persuader toward Right *reveals all arguments in favor of* Right, *and for $L \geq \bar{L}$ the persuader toward* Left *reveals all arguments in favor of* Left.

Figure 4 represents a continuous equilibrium for symmetric $f(L, R)$ and for $\pi_l = \pi_r$. The striped area represents reports used in equilibrium only by P_r and H, the dotted area represents reports used in equilibrium only by P_l and H, and the shaded square represents the set of reports that are used in equilibrium by all three types of expert, the *double ambiguity area.* As in the one-sided case, each type of persuaders reveals only a limited number of unfavorable arguments, and the decision maker disregards arguments that come in a short supply. Hence, the two main findings from section 3 still hold, namely, in the equilibrium

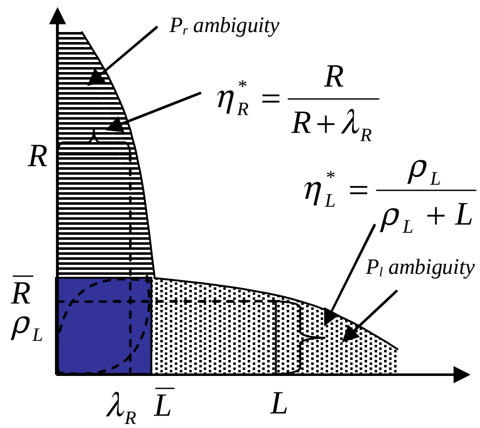

Figure 4: Two-sided bias.

the persuaders use two-sided messages, and, unless presented with balanced evidence, the decision maker bases her beliefs only on the arguments that favor the expert.

Ideally, each persuader would like to be perceived as unbiased or biased against his preferred alternative. However, in equilibrium when the persuader has sufficiently many favorable arguments, the benefit from disclosing that many favorable arguments exist dominates the loss of credibility from disclosing his potential bias.

The situation looks slightly different when the persuader's preferred alternative is of low quality. If the expert reveals few arguments in favor of *Right,* it is more likely that he has concealed a lot of arguments in favor of *Left.* Hence, if the expert reveals himself to be biased toward *Right,* the decision maker would rather choose *Left.* This is why the experts have no incentive to separate themselves when they have few favorable arguments. Therefore, for small λ and ρ the equilibrium resembles a babbling equilibrium. When the expert supports his recommendation with just few arguments, the decision maker ignores what the expert says.

In the example shown in Figure 4 when the decision maker sees a report from the double ambiguity area, she chooses the alternative supported by her prior belief. However, this need not to be true if the case is not symmetric. For example, for a symmetric $f(L, R)$, if $\pi_l > \pi_r$, then $\eta\left(\bar{L}, \bar{R}\right) > \frac{1}{2}$. In such a case, the decision maker expects that short reports are most likely to arise because the persuader toward *Left* conceals arguments unfavorable to him.

One may ask what would happen if the persuader could reveal his potential bias, for example, if a doctor could disclose who sponsored his research. The model says that if the issue is simple, the expert would not like to reveal this information. If the issue is complex, however, such disclosure is completely harmless. This finding sheds some light on why sometimes a party presenting a case, for example an author of an op-ed article, does disclose his or her conflict of interest. It also implies that mandatory disclosure of the conflict of interest might benefit the decision maker only when the issue on which she consults the expert turns out to be simple.

5 Confirmatory Bias

When confronted with mixed, ambiguous evidence, people tend to become even more extreme in their beliefs–this phenomenon is called *confirmatory bias* and have been documented in psychological literature. Darley and Gross (1983) asked seventy undergraduates to evaluate a nine-year-old girl's academic skills. Before the evaluation, the subjects were divided randomly into two groups: the first group watched a video of the girl playing in an affluent neighborhood and were told that girl's parents are college graduates with white-collar jobs; the second group watched a video of the girl playing in an impoverished neighborhood and were told that girl's parents are high school graduates with blue-collar jobs. Subsequently, half of each group were asked to evaluate the girl's academic skills, and, not surprisingly, the group that watched the girl playing in an affluent neighborhood provided a higher estimate. The remaining undergraduates were shown a second video of the same girl answering a sequence of questions with a mixed success, and then they were asked to evaluate the girl's academic skills. The estimates provided by these subjects were more polarized than the estimates provided by subjects who did not watch the second video.[11] In another experiment, Lord, Ross, and Lepper (1979) divided subjects into opponents and proponents of capital punishment based on self-reported attitudes. After presenting both groups with the same, mixed evidence on the deterrent effect of death penalty, the experimenters asked the subjects to evaluate the effect of this evidence on their beliefs. Compared to their initial beliefs, proponents reported greater belief in the deterrent effect of capital punishment, while opponents reported less belief in this deterrent effect.[12]

Confirmatory bias has been considered a departure from Bayesian updating and viewed as people's tendency to misinterpret ambiguous evidence as confirming their current hypotheses. Rabin and Schrag (1999) propose a non-Bayesian model of confirmatory bias and postulate accounting for confirmatory bias in economic models. To my knowledge, there is no rational model that could explain confirmatory bias.

[11] The estimates in the control groups were 4.29 and 3.9, respectively, and in the treatment groups were 4.67 and 3.71, respectively.

[12] For other studies, for example, see Plous (1991) and Jennings, Lepper, and Ross (1981).

My model fills this gap. When a possibly biased expert presents unbalanced evidence, a rational decision maker forms her belief using only on the arguments that favor what seems to be the expert's preferred alternative. Hence, when presented with subsequent mixed evidence, the decision maker might become even more enthusiastic about this alternative: she receives more favorable arguments, which causes her form a more positive belief about this alternative, and she also receives more unfavorable arguments, but these she disregards. Unless the subsequent evidence contains enough arguments against the initial position of the expert, which would suggest that the expert is honest, the confirmatory bias arises.

In the experiment of Darley and Gross (1983), the group that watched the girl playing in an affluent neighborhood received more arguments in favor of the girl, and could conclude that the experimenter might try to bias the subject's estimates upward. My model predicts then that they should base their estimates only on the favorable arguments. As a result, the treatment group, which was exposed to more favorable arguments than the control group, should have a more positive view of the girl's skills. Similarly, the subjects in the treatment group who watched the video of the girl playing in the impoverished neighborhood should have a more negative view of the girl's skills than the subjects in the control group.

A skeptical reader might argue that there is no reason for the experimenter to be biased. However, subjects in the laboratory experiments do not necessarily know or understand the objective of each experiment; and hence, they might expect the experimenter to be biased or to be under influence of biased information. Moreover, people's behavior in the laboratory might simply mimic their behavior from real life, where attempts to manipulate information are abundant.

My model also explains why people with different priors might come to different conclusions even if they end up having the same information. If the different priors mean that people have different prior distributions over R and L, the decision maker who leans initially toward the option favored by the expert will form a higher belief than the decision maker with a less favorable prior. For each number of favorable arguments, the number of unfavorable arguments that the decision maker expects the expert to conceal is higher for the second type of the decision maker. My model does not explain, however, why the beliefs of people with different priors should become polarized after a sequence of mixed evidence,[13] which is the behavior reported by Lord, Ross, and Lepper (1979). As Rabin and Schrag (1999) mention, however, the experiments reported by Lord, Ross, and Lepper (1979) might represent a phenomenon different than confirmatory bias. People might have different priors specifically because they might have a tendency to read evidence in a biased way, and ob-

[13]Such a phenomenon could in principle happen in my model, but only for specific prior beliefs and realization of the signals.

serving further polarization by groups who already differ may reflect underlying differences in interpretation of evidence that would appear irrespective of people' current beliefs. Such a behavior is clearly non-Bayesian.

6 Comparative statics

In this section, I analyze how the equilibrium is affected by the parameters of the model, such as the probabilities of different types of the expert and the prior distribution of arguments.

6.1 Varying the probability of facing the persuader

This section analyzes how changes in the probability of facing the persuader affect the players' utilities and, more generally, the whole equilibrium. First, I look at what happens when the fraction of honest experts becomes small and what happens when the expert is honest with probability of almost 1. Second, I analyze how the probability of facing a particular type of the persuader impacts the bias of his reports and the probability of persuading the decision maker.

Since there are three types of expert, it is necessary to specify how the remaining probabilities change when the probability of facing the expert of type i changes. In the second part of Proposition 5, I vary the probability of P_r and keep constant the conditional probability of facing the honest type, given that the expert is not P_r. In such a case the shape of the ambiguity area for P_l remains the same.

Proposition 5 *In every continuous equilibrium,* $\lim_{\pi_H \to 1} \lambda_R = 0$, $\lim_{\pi_H \to 1} \rho_L = 0$, *and* $\lim_{\pi_H \to 0} \lambda_R = \bar{\lambda}_R$, $\lim_{\pi_H \to 0} \rho_L = \bar{\rho}_L$, *where* $\bar{\lambda}_R$ *and* $\bar{\rho}_L$ *are such that* $\frac{R}{R+\bar{\lambda}_R} = E\left[q_R|R\right]$ *and* $\frac{\bar{\rho}_L}{\bar{\rho}_L+L} = E\left[q_R|L\right]$. *With* $\frac{\pi_H}{1-\pi_r}$ *kept constant, as the probability of facing* P_r *decreases,*

i. the reports of P_r *become more extreme;*

ii. the utility of P_r *increases;*

iii. the utility of P_l *decreases;*

iv. the expected utility of the decision maker increases.

Proposition 5 says that as the probability of facing the honest expert increases, the ambiguity areas for both persuaders disappear, and the equilibrium converges to a fully revealing equilibrium. On the other hand, as the probability of facing the honest expert goes down, λ_R converges to $\bar{\lambda}_R < \infty$ and ρ_L converges to $\bar{\rho}_L < \infty$. This means that the ambiguity areas are always strict subsets of the report space and the equilibrium never becomes a pure babbling equilibrium. The assumption that arguments are verifiable prevents equilibria from becoming completely uninformative.

It is worth noting that even when the expert is a persuader with probability 1, $\pi_H = 0$, revealing unfavorable information can be a part of an equilibrium. In the one-sided case when $\pi_H = 0$, it is possible to construct an equilibrium in which the persuader suppresses the unfavorable information completely, but the decision maker is still skeptical about all reports with a small number of unfavorable arguments. In particular, since she may attach the same belief to all such reports, it does not hurt the persuader to reveal some unfavorable arguments. In a two-sided case the situation is even more interesting. Even if $\pi_H = 0$, there cannot be an equilibrium with full suppression of unfavorable information. Each persuader would like his type to be misperceived by the decision maker so that her skepticism works in his favor; therefore, when the persuaders receive very few favorable arguments, they must pool. Pooling can happen only if none of the persuaders completely suppress unfavorable information.

This finding reveals an important difference between the current model and Shin (1994a) and Shin (1994b). These papers analyze a similar model but with common knowledge about the preferences of the expert, and they focus solely on the equilibrium in which only the favorable information is revealed. My paper shows that when the state space is continuous, such an equilibrium is not robust to the uncertainty about the expert's type. In section 6.3 I discuss how my results would extend to a case with a discrete state space. When the state space is discrete, it is possible that the persuader completely suppresses the unfavorable arguments, but this will be an equilibrium only under very specific distributional assumptions.

Figure 5 shows how the equilibrium changes as π_r decreases. π_1 and $\pi_0 > \pi_1$ are two different probabilities of facing P_r. The thick curves represent the initial equilibrium in which $\pi_r = \pi_0 = \pi_l$. Since the conditional probability of facing H is kept constant, the shape of the ambiguity area for P_l remains unchanged as π_r changes. As $\pi_r \to 0$, the ambiguity area for P_r becomes smaller, as represented by the thinner curve. The shaded region shows the double ambiguity area for $\pi_r = \pi_1$.

When π_r decreases, the reports of P_r become more extreme. The decision maker becomes less skeptical about the biased reports since they are less likely to come from the persuader; hence, she attaches a higher belief to them. P_r will not send more arguments in favor of *Left,* because this would prove that the quality of *Right* is lower.

From Figure 5 one can see that when the decision maker faces P_r, she chooses *Right* more often when π_r is low (for π_0, a decision maker with θ chooses *Right* whenever $R \geq R_\theta^{\pi_0}$, and

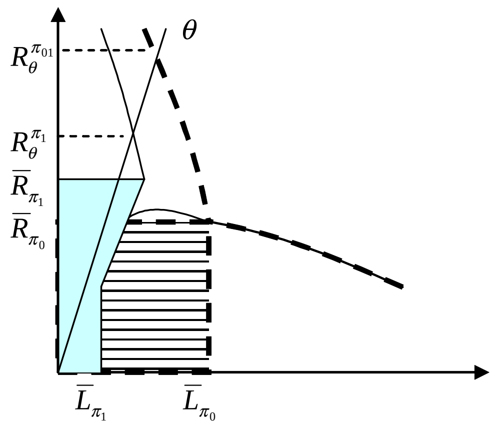

Figure 5: The effect of decreasing π_r, while keeping $\frac{\pi_H}{1-\pi_r}$ constant.

for π_1, whenever $R \geq R_\theta^{\pi_1}$). Therefore, the utility of the persuader toward *Right* increases. If the expert happens to be P_l, the decision maker chooses *Right* more often, which decreases the utility of the persuader toward *Left.* However, the utility of the decision maker increases because she is more likely to face the honest expert.

Proposition 5 implies that a financial adviser biased toward a stock that is unpopular among other advisers is better at persuading investors to buy that stock, while a financial adviser biased toward a popular stock is unlikely to successfully promote it.

6.2 Varying the familiarity of the problem

In this section, I analyze how the prior distribution of arguments affects the utility of the decision maker. I focus on the case in which the potential bias of the persuader is known, i.e., when the expert can be either P_r or H.

With the distribution of quality q_R held constant, the distribution of the total number of arguments, $N \equiv L + R$, reflects the decision maker's uncertainty about the choice problem. It describes how the total number of arguments varies from situation to situation for the same decision problem. For example, in each election campaign a different number of issues is important, which can be represented in the model by a relatively dispersed prior belief over N. Other choice problems are likely to be characterized by roughly the same number of arguments every time the decision maker faces them, such us choosing an investment option or buying a car; this is captured by a distribution of N concentrated around the mean. Alternatively, the prior distribution of N may describe the decision maker's knowledge about the problem. An investor with a dispersed distribution of N knows little about the nature of the problem, while an experienced or educated investor is likely to have a concentrated distribution of N.

To isolate the effect of changing the distribution of N while keeping the distribution of the quality of the alternatives unchanged, I reformulate the problem in terms of $(q_R, N) = \left(\frac{R}{R+L}, L+R\right)$, and assume that q_R is uniformly distributed and independent of N. This implies that the joint density of q_R and N is equal to the density of N. Let $G(N;z)$ be the family of distributions with the following properties:

1. the expected value of N is independent of z, $\bar{N} \equiv \int_0^\infty N dG(N;z)$;

2. for all z and for all N we have $G_z(N;z) < 0$ if $N < \bar{N}$ and $G_z(N;z) > 0$ if $N > \bar{N}$;

3. as $z \to \infty$, $G(\cdot)$ becomes degenerate at $\bar{N}$.

Assume that N is distributed according to $G(N;z)$. The second property says that the higher z, the more centered around the mean the distribution is.

Proposition 6 *For every preference type of the decision maker θ and every $\pi > 0$, if $z_1 > z_2$, then the decision maker's utility is higher for $G(N;z_1)$ than for $G(N;z_2)$. As $z \to \infty$, there is full revelation of information.*

Proposition 6 says the lower the uncertainty about N, the better-off the decision maker. When the decision maker knows more about how many arguments are available to the expert, she can more easily infer his information: when she receives a report, she can estimate rather precisely how many arguments have been concealed from her. Given that the dispersion of N can represent the decision maker's familiarity with the choice problem or general uncertainty about it, this proposition implies that the decision maker is better-off in familiar situations with constant complexity.

6.3 Robustness

6.3.1 Selection of continuous equilibria

For most of the paper, I focused on the equilibria with the continuous belief function because all equilibria share the same qualitative features, and the continuous equilibria seem more appealing. Proposition 7 establishes another attractive feature of continuous equilibria.[14]

Proposition 7 *The ex-ante utility of the decision maker is the highest in continuous equilibria.*

[14] The previous version of this paper contained a proposition that if we perturb the game by adding a small fixed cost of concealing information, its unique equilibrium converges to a continuous equilibrium. Results upon request.

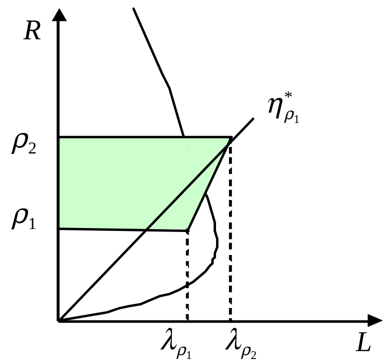

Figure 6: An example of a discontinuous equilibrium.

In continuous equilibria, the persuader induces a different belief for each R, while this is not true in discontinuous equilibria. We can find $R_1 > R_2$ such that some types of the decision maker, when facing the persuader, choose *Right* when $R = R_1$ and *Left* when $R = R_2$ in a continuous equilibrium, but choose the same alternative in both cases in some discontinuous equilibrium. This suggests that the decision maker is worse-off in this discontinuous equilibrium if she faces the persuader. She may be better-off when she faces the honest expert if the ambiguity area is smaller, in which case the honest expert can signal higher qualities of *Right* than he could in continuous equilibria. It turns out that the loss due to worse decision making when facing the persuader always outweighs the benefit from better decision making when facing the honest expert. The reason for this is that the affected decision maker has θ high enough that the probability of facing the honest expert and a state being $q_R > \theta$ – that is, the probability of improving the decision – is relatively low.

I want to shed more light, however, on what discontinuous equilibria look like. Figure 6 shows a representative equilibrium for the case with one type of the persuader. In this equilibrium the belief induced when only few arguments in favor of *Left* are revealed is strictly increasing for $\rho \in (0, \rho_1)$ and $\rho \in (\rho_2, \infty)$ and is constant for $\rho \in (\rho_1, \rho_2)$. The curve represents λ_ρ for $\rho \in (0, \rho_1)$ and $\rho \in (\rho_2, \infty)$. The shaded trapezoid represents reports that generate the same belief $\eta^*_{\rho_1}$. Additionally, all reports lying in the rectangle that completes the trapezoid generate either the same belief as the trapezoid, or are sent only by H.

6.3.2 Utility function

In this model the decision maker cares only about the quality of each alternative; the total number of arguments does not enter her utility function. This reflects the idea that the set of all arguments carries *perfect* information about the state of nature, and the number of arguments measures only the complexity of the problem. For example, computers have more features than memory cards, and buying a house requires considering more aspects than buying a couch. The decision maker is uncertain about this complexity, but the complexity itself does not affect her preferences. There is no reason why consumers should require stronger proof when buying a more complex product.

The qualitative results of this paper still hold if the decision maker's utility depends on the quality of the chosen alternative, but the expected value of the quality depends on the total number of arguments as well as their proportion. In equilibrium, the expected quality given the expert's report must be constant for some λ; otherwise the expert would use only the reports generating the highest expectation. More generally, for any utility specification the probability of *Right* being chosen must be constant for any small number of unfavorable arguments. Hence, in all these models the decision maker bases her decision only on the number of favorable arguments, and the expert must use two-sided messages.

6.3.3 Continuity of arguments

The intuition for Proposition 2 makes it clear that the assumption of the continuity of arguments is not completely innocuous. If arguments are discrete, revealing one unfavorable argument proves that the quality of *Right* is significantly lower than 1, and this may outweigh the benefit from gaining credibility. However, complete suppression of unfavorable arguments may be a part of an equilibrium only if the probability of a state with no arguments in favor of *Left* is relatively high compared to the probability of the expert being a persuader and having at least one such argument. If the distribution of arguments is not very skewed toward states with arguments only in favor of *Right*, and the probability that the expert is a persuader is high enough, then the main conclusion of Proposition 2 would hold in a discrete model.

6.3.4 Benevolent expert

In this model the honest expert reveals all of the arguments. Alternatively, the honest expert may want to maximize the utility of the decision maker, i.e., he may be benevolent. One can easily see, however, that any continuous equilibrium of the original game is still an equilibrium of a game with a benevolent type of expert, and the benevolent expert behaves like an honest expert. To see this, note that if the state of nature lies in the revealing area, the benevolent expert cannot do better than to report fully, because in this way he induces the correct belief. If the state of nature lies in the ambiguity area, however, the benevolent expert would like to induce a higher belief than the one induced in equilibrium, but there is no feasible report that can achieve this; therefore, full reporting is once again optimal.

When the expert is benevolent, however, there are more equilibria. In particular, it is no longer true that all reports must be used. As a result, by appropriately choosing the off-equilibria beliefs, one can support many implausible equilibria. Hence, assuming that the expert is honest is similar to using a refinement which requires off-equilibrium reports to be taken at face value.

7 Conclusion

This paper proposes a model of communication that captures many economically relevant situations. In the model messages available to the expert are verifiable, but the expert cannot prove whether he has disclosed all information. As a result, unravelling fails. The decision maker takes balanced reports at face value, but is skeptical about the unbalanced ones. In such a case, she chooses the alternative favored by the persuader only if the expert can provide her with sufficiently many favorable arguments. If the decision maker does not know the bias of the expert, she ignores his recommendation unless he can reveal sufficiently many arguments. She also does better if she knows more about the complexity of the problem, and is more likely to be persuaded in complex situations.

This paper provides a game theoretic foundation for the use of two-sided messages, and can account for confirmatory bias in a rational setting. It also contributes to the research on whether mandatory disclosure improves the decision maker's welfare by demonstrating that the answer is positive. Disclosure is always beneficial and has the greatest value when the decision maker is unfamiliar with the problem and the complexity of the problem is volatile. However, mandatory disclosure may be difficult to implement if it is difficult to prove that the expert was informed in the first place. This model suggests that a policy of educating the decision maker about which arguments, facts, or characteristics are relevant will improve the decision maker's welfare.

To prove the relevance of my results, in accordance with my model I should also mention some limitations of it. First, the model assumes that the decision maker is uninformed, but it would clearly be interesting to analyze the case in which the decision maker has some prior information. Second, the model explains the use of two-sided messages, but it predicts that two-sided messages are not more persuasive. Although I do not think that two-sided messages are always more persuasive, experiments on mass communication indicate that two-sided arguments might be more effective.[15] Slight modification of the model, such as

[15] See Hovland, Lumsdaine, and Sheffield (1949) and Lumsdaine and Janis (1953).

the introduction of a small fraction of decision makers who take the experts' reports at face value, could generate this effect; this, however, is left for future research.

The model presented in this paper provides a good starting point for analyzing communication in more elaborate settings. Because of the argument structure, the differences in information that players have is easily defined, which makes this model especially well-suited for analyzing two-sided communication. Moreover, since disclosing arguments requires time, the model has some natural timing structure built in, which means that it can be applied to dynamic communication.

A Appendix

Proof of Proposition 1

Pick a small λ_0 and consider a set $X(\rho_0) = \{(\lambda, \rho) : \lambda \in (0, \lambda_0) \text{ and } \rho \in (\rho_0, \infty)\}$. All reports in this set have very high proportion of arguments in favor of *Right,* and as $\rho_0 \to \infty$, $E\left[\frac{R}{R+L} | (L, R) \in X(\rho_0)\right] \to 1$. For the decision maker to form a belief equal to the prior, $\eta(\lambda, \rho) = E\left[\frac{R}{R+L}\right] \equiv \bar{\eta}$, after all reports in $X(\rho_0)$, there must exist a set of states $Y(\rho_0)$ in which quality of *Right* is low, such that the persuader P_r sends reports from $X(\rho_0)$ when $(L, R) \in Y(\rho_0)$. Below, I show that for a sufficiently high ρ_0, it is impossible to find a set of states $Y(\rho_0)$ in which reports from $X(\rho_0)$ are feasible such that if P_r sends reports from $X(\rho_0)$ whenever $(L, R) \in Y(\rho_0)$ the belief that the decision maker forms is at least $\bar{\eta}$.

Take a set $Y(\rho_0) = \left\{(L, R) : R \geq \rho_0, L \geq \frac{1-\bar{\eta}}{\bar{\eta}} R\right\}$. This set contains all states, and only these states, in which the quality of *Right* is lower than $\bar{\eta}$. If we are unable to generate a belief equal to or below $\bar{\eta}$ for such $Y(\rho_0)$, then we are unable to do this for any other feasible set.

The expected quality of *Right* given that the report is in X is

$$\begin{aligned} E\left[\frac{R}{R+L} | (\lambda, \rho) \in X\right] &= \Pr(H | (\lambda, \rho) \in X) E\left[\frac{R}{R+L} | (L, R) \in X\right] \\ &+ (1 - \Pr(H | (\lambda, \rho) \in X)) E\left[\frac{R}{R+L} | (L, R) \in Y\right]. \end{aligned}$$

Let $f^R(\cdot)$ denote the marginal p.d.f. of R. We have

$$\begin{aligned}
\lim_{\rho_0\to\infty}\frac{\int_{\rho_0}^{\infty}\int_{\frac{1-\bar{\eta}}{\bar{\eta}}R}^{\infty}\frac{R}{R+L}f^L(L|R)\,f^R(R)\,dLdR}{\int_{\rho_0}^{\infty}\int_{0}^{\lambda_0}f^L(L|R)\,f^R(R)\,dLdR} &\leq \lim_{\rho_0\to\infty}\frac{\int_{\rho_0}^{\infty}\int_{\frac{1-\bar{\eta}}{\bar{\eta}}R}^{\infty}f^L(L|R)\,f^R(R)\,dLdR}{\int_{\rho_0}^{\infty}\int_{0}^{\lambda_0}f^L(L|R)\,f^R(R)\,dLdR} =^H= \\
&= \lim_{\rho_0\to\infty}\frac{\int_{\frac{1-\bar{\eta}}{\bar{\eta}}\rho_0}^{\infty}f(L|\rho_0)\,dLf^R(\rho_0)}{\int_{0}^{\lambda_0}f^L(L|\rho_0)\,dLf^R(\rho_0)} \\
&= \lim_{\rho_0\to\infty}\frac{1-F^L\left(\frac{1-\bar{\eta}}{\bar{\eta}}\rho_0|\rho_0\right)}{F^L(\lambda_0|\rho_0)} = 0,
\end{aligned}$$

where $=^H=$ denotes that l'Hospital rule was applied and the last equality comes from the regularity conditions 2. Since $\frac{dF^L(L|R)}{dR}\geq 0$, we have $\lim_{\rho_0\to\infty}F^L(\lambda_0|\rho_0)>0$, while $\lim_{\rho_0\to\infty}F^L\left(\frac{1-\bar{\eta}}{\bar{\eta}}\rho_0|\rho_0\right)=1$.

Therefore, we have

$$\begin{aligned}
&\lim_{\rho_0\to\infty}\Pr\left(H|\left(\lambda,\rho\right)\in X\right) \\
=&\lim_{\rho_0\to\infty}\frac{(1-\pi)\int_{\rho_0}^{\infty}\int_{0}^{\lambda_0}f^L(L|R)\,f^R(R)\,dLdR}{(1-\pi)\int_{\rho_0}^{\infty}\int_{0}^{\lambda_0}f^L(L|R)\,f^R(R)\,dLdR+\pi\int_{\rho_0}^{\infty}\int_{\frac{1-\bar{\eta}}{\bar{\eta}}R}^{\infty}f^L(L|R)\,f^R(R)\,dLdR}=1,
\end{aligned}$$

hence

$$\lim_{\rho_0\to\infty}E\left[\frac{R}{R+L}|\left(\lambda,\rho\right)\in X\right]=1>\bar{\eta},$$

which completes the proof. ∎

Proof of Proposition 2

Let $\sigma_i(L,R)$ denote the set of all reports that lie in the support of the strategy of an expert of type i, that is, $\sigma_i(L,R)=\{(\lambda,\rho):m_i((\lambda,\rho)\,|\,(L,R))>0\}$. Recall that $\eta(\lambda,\rho)=E\left[\frac{R}{R+L}|\left(\lambda,\rho\right)\right]$ and $Z(\lambda,\rho)=((L,R):L\geq\lambda\text{ and }R\geq\rho)$ is a set of all states in which a report (λ,ρ) is feasible.

Step 1 The presence of the honest expert, H, assures that there are no off-equilibrium beliefs: whenever $(\lambda,\rho)\notin\sigma_{P_r}(L,R)$ for all (L,R) (that is, whenever (λ,ρ) is never sent by P_r) then $\eta(\lambda,\rho)=\frac{\rho}{\lambda+\rho}$.

Step 2 $\eta(0,\rho)<1$ for all ρ.

By step 1, $\eta(\lambda,\rho)<1$ for all $\lambda>0$. Assume there exist ρ such that $\eta(0,\rho)=1$. Then for all $(L,R)\in Z(0,\rho)$, P_r is able to generate a belief equal to 1, which would violate Bayes' rule.

Step 3 For all ρ, there exists $l_\rho>0$ such that $\eta(\lambda,\rho)=\eta(0,\rho)$ for all $\lambda<l_\rho$.

Assume not. This means that for some ρ and for all $\varepsilon > 0$, we can find $\lambda_0 < \varepsilon$ such that $\eta(\lambda_0, \rho) \neq \eta(0, \rho)$. Assume first that $\eta(\lambda_0, \rho) < \eta(0, \rho)$. By Step 2 $\eta(\lambda_0, \rho) < \eta(0, \rho) < 1$. But then P_r always prefers to send $(0, \rho)$ instead of (λ_0, ρ), which implies that (λ_0, ρ) is sent only by H. Step 1 implies that $\eta(\rho, \lambda_0) = \frac{\rho}{\lambda_0+\rho}$. But $\eta(\rho, \lambda_0) = \frac{\rho}{\lambda_0+\rho}$ becomes arbitrarily close to 1 as $\varepsilon \to 0$. This contradicts the assumption that $\eta(\lambda_0, \rho) < \eta(0, \rho) < 1$.

Assume then that for all $\varepsilon > 0$ we can find $\lambda_0 < \varepsilon$ such that $\eta(\lambda_0, \rho) > \eta(0, \rho)$. Then P_r may send $(0, \rho)$ only if $L = 0$ because for any $L > 0$ he would find $\lambda_0 \leq L$ and send (λ_0, ρ) instead. But then $\eta(0, \rho) = 1$, which contradicts Step 2.

Define $\eta^*_\rho \equiv \eta(0, \rho)$.

Step 4 η^*_ρ is weakly increasing in ρ.

Assume that there exist $\rho_2 > \rho_1$ such that $\eta^*_{\rho_2} < \eta^*_{\rho_1}$. Then P_r never sends $(0, \rho_2)$, which by Step 1 implies that $\eta^*_{\rho_2} = 1$, which contradicts Step 2.

Step 1 and Step 3 imply that in equilibrium P_r does not always completely suppress unfavorable information; otherwise $\eta(\lambda, \rho)$ would be a strictly decreasing function of λ.

Step 5 $\eta^*_\rho \geq \frac{\rho}{\rho+\lambda_\rho}$

Assume $\frac{\rho}{\rho+\lambda_\rho} > \eta^*_\rho$. Then for all $\varepsilon > 0$, we can find $\lambda_0 \in (\lambda_\rho, \lambda_\rho + \varepsilon)$ such that $\eta(\lambda_0, \rho) \neq \eta^*_\rho$. Assume first that $\eta(\lambda_0, \rho) < \eta^*_\rho$ Then only H sends (λ_0, ρ), and therefore, $\eta(\lambda_0, \rho) = \frac{\rho}{\rho+\lambda_0} > \eta^*_\rho$, a contradiction. Assume now that $\eta(\lambda_0, \rho) > \eta^*_\rho$. But then (λ_0, ρ) is more attractive to P_r than any report of a form $(\lambda \leq \lambda_\rho, \rho)$. Therefore, P_r will send (λ_ρ, ρ) only if (L, R) is such that $\frac{R}{R+L} \geq \frac{\rho}{\rho+\lambda_\rho} > \eta^*_\rho$, which contradicts $\eta(\lambda_\rho, \rho) = \eta^*_\rho$.

Step 6 Let $\hat{\rho} = \max\left\{\rho : \eta^*_\rho = \hat{\eta}\right\}$ for some $\hat{\eta}$. Then $\lambda_{\hat{\rho}}$ is such that $\frac{\hat{\rho}}{\hat{\rho}+\lambda_{\hat{\rho}}} = \hat{\eta}$.

By Step 5 we have $\frac{\hat{\rho}}{\hat{\rho}+\lambda_{\hat{\rho}}} \leq \hat{\eta}$. Assume then $\frac{\hat{\rho}}{\hat{\rho}+\lambda_{\hat{\rho}}} < \hat{\eta}$. To generate a belief $\hat{\eta}$ for a report $(\lambda_{\hat{\rho}}, \hat{\rho})$ it must be that there exists $(L_1, R_1) \in Z(\lambda_{\hat{\rho}}, \hat{\rho})$ such that $\frac{R_1}{R_1+L_1} \geq \hat{\eta}$ and $\sigma_{P_r}(L_1, R_1) = (\lambda_{\hat{\rho}}, \hat{\rho})$. But that implies that $R_1 > \hat{\rho}$, and by definition of $\hat{\rho}$ we know $\eta^*_{R_1} > \hat{\eta}$, therefore, P_r would prefer to send $(0, R_1)$ to sending $(\lambda_{\hat{\rho}}, \hat{\rho})$. A contradiction. ■

Proof of Proposition 3

I characterize all equilibria in which the belief function $\eta(\lambda, \rho)$ is continuous in λ and ρ. First, continuity requires that $\eta^*_\rho = \frac{\rho}{\lambda_\rho+\rho}$; therefore, there exists a unique λ_ρ for a given η^*_ρ. This, together with continuity, implies that there is no $\lambda > \lambda_\rho$ such that $\eta(\lambda, \rho) \geq \eta^*_\rho$, which implies that only H sends reports of a form $(\lambda > \lambda_\rho, \rho)$. Steps 1-4 from the proof of Proposition 2 together with $\eta^*_\rho = \frac{\rho}{\lambda_\rho+\rho}$ imply that all equilibria have a shape like the one in Figure 6. The solid curve represents λ_ρ for each ρ. Given that $\eta^*_\rho = \frac{\rho}{\lambda_\rho+\rho}$ the equilibrium depicted above is characterized by a strictly increasing η^*_ρ but for $\rho \in (\rho_1, \rho_2)$. Now, I show that $\eta^*_{\rho_1} < \eta^*_{\rho_2}$, hence this equilibrium is not continuous. Therefore in a continuous equilibrium η^*_ρ must be strictly increasing.

First, we can find ε such that η^*_ρ is strictly increasing in $\rho^{'}$ for all $\rho' \in (\rho_1, \rho_1 + \varepsilon)$ and $\rho' \in (\rho_2 - \varepsilon, \rho_2)$. If we could not, the equilibrium would not be continuous. For $R = \rho'$ the persuader reveals R and sends some $\lambda < \lambda_{\rho'}$; therefore, the belief function for these ρ's is defined by

$$\begin{aligned}\eta^*_{\rho'} &= \Pr\left(H|\rho\prime, \lambda \leq \lambda_{\rho'}\right) E\left[q_R | R = \rho\prime, L \leq \lambda_{\rho'}\right] \\ &\quad + \Pr\left(P_r|\rho\prime, \lambda \leq \lambda_{\rho\prime}\right) E\left[q_R | R = \rho\prime\right].\end{aligned}$$

where $\lambda_{\rho'} = \frac{1-\eta_{\rho'}}{\eta_{\rho'}}\rho'$. If $\eta(\rho, \lambda)$ is continuous then $\lim_{\varepsilon \to 0} \eta^*_{\rho_2 + \varepsilon} = \lim_{\varepsilon \to 0} \eta^*_{\rho_1 - \varepsilon}$.

Let $g(\rho)$ be η^*_ρ that satisfied the above equation for each ρ. This equation can be rewritten as:

$$\begin{aligned}Y \equiv g(\rho)\left(F^L\left(\frac{1-g(\rho)}{g(\rho)}\rho | R = \rho\right)(1-\pi) + \pi\right) + \\ -(1-\pi)\int_0^{\frac{1-g(\rho)}{g(\rho)}\rho} \frac{\rho f^L(L|R=\rho)}{\rho + L} dL \quad -\pi \int_0^\infty \frac{\rho}{\rho+L} f^L(L|R=\rho)\, dL = \quad 0.\end{aligned} \tag{5}$$

Clearly, $g(\rho_2) = \lim_{\varepsilon \to 0} \eta^*_{\rho_2+\varepsilon} = \eta^*_{\rho_2}$ and $g(\rho_1) = \lim_{\varepsilon \to 0} \eta^*_{\rho_1 - \varepsilon} = \eta^*_{\rho_1}$. I now show that $\frac{dg(\rho)}{d\rho} > 0$, which implies that $g(\rho_2) > g(\rho_1)$, which in turn means that $\eta^*_{\rho_2} > \eta^*_{\rho_1}$, which is a contradiction.

By the implicit function theorem $\frac{dg(\rho)}{d\rho} = -\frac{\frac{\partial Y}{\partial \rho}}{\frac{\partial Y}{\partial g}}$, and

$$\frac{\partial Y}{\partial g} = F^L\left(\frac{1-g}{g}\rho | R = \rho\right)(1-\pi) + \pi > 0,$$

$$\begin{aligned}\frac{\partial Y}{\partial \rho} &= -(1-\pi)\int_0^{\frac{1-g}{g}\rho} \frac{L f^L(L|R=\rho)}{(\rho+L)^2} dL - \pi \int_0^\infty \frac{L f^L(L|R=\rho)}{(\rho+L)^2} dL \\ &\quad -(1-\pi)\int_0^{\frac{1-g}{g}\rho} \frac{\rho F^L_R(L|R=\rho)}{(\rho+L)^2} dL - \pi \int_0^\infty \frac{\rho F^L_R(L|R=\rho)}{(\rho+L)^2} dL,\end{aligned}$$

where the expression above was obtained by applying first integration by parts to the formula for Y, taking the derivative and applying integration by parts again. Using the regularity condition (2) we get $\frac{dY}{d\rho} < 0$, which in turn implies that $\frac{dg(\rho)}{d\rho} > 0$.

When η^*_ρ is strictly increasing, P_r sends always $\rho = R$. Given this, I sometimes use η^*_R instead of η^*_ρ.

It remains to show that for each R, η^*_R described by equation (3) exists and is unique.

Equation (3) can be rewritten as:

$$\eta_R^* = \frac{\pi \int_0^\infty \frac{R}{R+L} f^L(L|R)\, dL + (1-\pi) \int_0^{\frac{1-\eta_R^*}{\eta_R^*}R} \frac{R}{R+L} f^L(L|R)\, dL}{\pi + (1-\pi) F^L\left(\frac{1-\eta_R^*}{\eta_R^*} R|R\right)}. \tag{6}$$

The left hand side (LHS) goes from 0 to 1. The right hand side (RHS) is continuous, and we have $RHS(\eta_R^* \to 0) \to \bar{\eta}_R > 0$ and $RHS(\eta_R^* = 1) = \bar{\eta}_R < 1$, where $\bar{\eta}_R = \int_0^\infty \frac{R}{R+L} f^L(L|R)\, dL$; therefore, the solution to equation (6) exists.

The LHS is strictly increasing. If we differentiate the RHS with respect to η_R^*, we obtain

$$\begin{aligned} \frac{dRHS}{d\eta_R^*} \;=\; & \frac{(1-\pi) f^L\left(\frac{1-\eta_R^*}{\eta_R^*} R|R\right)}{\left(\pi + (1-\pi) F^L\left(\frac{1-\eta_R^*}{\eta_R^*} R|R\right)\right)^2} \frac{-1}{(\eta_R^*)^2} R \cdot \\ & \cdot \left(\pi\left(\eta_R^* - \bar{\eta}_R\right) + (1-\pi)\left(\eta_R^* F^L\left(\frac{1-\eta_R^*}{\eta_R^*} R|R\right) - \int_0^{\frac{1-\eta_R^*}{\eta_R^*}R} \frac{R f^L(L|R)}{R+L} dL\right)\right). \end{aligned}$$

Evaluating it at η_R^* that satisfies equation (6) we obtain

$$\frac{dRHS}{d\eta_R^*} = \frac{(1-\pi) f^L\left(\frac{1-\eta_R^*}{\eta_R^*} R|R\right)}{\left(\pi + (1-\pi) F^L\left(\frac{1-\eta_R^*}{\eta_R^*} R|R\right)\right)} (\eta^* - \eta^*) = 0.$$

This implies that the solution to equation (6) is unique. ■

Proof of Proposition 4

Step 1 If for some (λ, ρ) there exists no (L, R) such that $(\lambda, \rho) \in \sigma_{P_r}(L, R)$ or $(\lambda, \rho) \in \sigma_{P_l}(L, R)$ then $\eta(\lambda, \rho) = \frac{\lambda}{\lambda+\rho}$.

This follows directly from the presence of H type. This finding will be used extensively in the proof.

Step 2 $\eta(0, \rho) < 1$ for all ρ and $\eta(\lambda, 0) > 0$ for all λ.

The proof of the first part is identical to Step 2 of the proof of Proposition 2. By a similar argument we get $\eta(\lambda, 0) > 0$.

Step 3 i) For all ρ, there exists $\lambda_\rho > 0$ such that $\eta(\lambda, \rho) = \eta(0, \rho)$ for all $\lambda < \lambda_\rho$.

ii) For all λ, there exists $\rho_\lambda > 0$ such that $\eta(\lambda, \rho) = \eta(\lambda, 0)$ for all $\rho < \rho_\lambda$

Assume that part (i) does not hold, which means that there exists $\lambda_0 > 0$ such that $\eta(\lambda, \rho)$ is strictly increasing or decreasing in λ for $\lambda < \lambda_0$. If $\eta(\lambda, \rho)$ is decreasing in λ then P_r always prefers to send $(0, \rho)$ instead of any (λ, ρ) with $\lambda < \lambda_0$ and P_l prefers to send as high λ as possible for $\lambda \in (0, \lambda_0)$. Take (ε, ρ) with $\varepsilon < \lambda_0$. P_l may send (ε, ρ) only if $L = \varepsilon$ and $R \geq \rho$, therefore it must be that $\eta(\varepsilon, \rho) \geq \frac{\rho}{\rho+\varepsilon}$. But $\eta(\varepsilon, \rho) \geq \frac{\rho}{\rho+\varepsilon} \rightarrow_{\varepsilon \rightarrow 0} 1 > \eta(0, \rho)$, which contradicts the continuity assumption.

Assume now that (λ, ρ) is strictly increasing in λ for $\lambda \leq \lambda_0$. Then P_l never sends (ε, ρ) for small ε : he would prefer to send $(0, \rho)$ instead. P_r may send (ε, ρ) only if $L = \lambda_\rho$ and $R \geq \rho$. But then $\eta(\varepsilon, \rho) \geq \frac{\rho}{\rho+\varepsilon} \rightarrow_{\varepsilon \rightarrow 0} 1$, which again contradicts the assumption of continuity of $\eta(\lambda, \rho)$. A similar argument holds for part ii).

Define $\eta^*_\rho \equiv \eta(0, \rho)$ and $\eta^*_\lambda \equiv \eta(\lambda, 0)$.

Step 4 i) P_r sends only reports of the form $(\lambda, \rho) : \lambda \leq \lambda_\rho$,

ii) P_l sends only reports of the form $(\lambda, \rho) : \rho \leq \rho_\lambda$.

Assume that P_r sends (λ_0, ρ_0), where $\lambda_0 > \lambda_{\rho_0}$. This means that $\eta(\lambda_0, \rho_0) \geq \eta^*_{\rho_0}$. This in turn means that for some $\lambda \in (\lambda_{\rho_0}, \lambda_0)$, the belief function $\eta(\lambda, \rho_0)$ is strictly increasing in λ but $\eta(\lambda, \rho_0) < \eta^*_{\rho_0}$. But these reports would be sent by H only, which means that $\eta(\lambda, \rho_0) = \frac{\rho_0}{\lambda+\rho_0}$. This contradicts that $\eta(\lambda, \rho_0)$ was increasing in λ. An analogous proof holds for part (ii).

By the same arguments we have $\eta^*_\rho > \eta(\lambda, \rho)$ for all $\lambda > \lambda_\rho$ and $\eta^*_\lambda < \eta(\lambda, \rho)$ for all $\rho > \rho_\lambda$.

Step 5 Previous steps imply that there exist $\bar{\rho} > 0$ and $\bar{\lambda} > 0$ such that $\eta^*_\rho = \eta^*_\lambda \equiv \eta^*$ for all $\rho \leq \bar{\rho}$ and for all $\lambda \leq \bar{\lambda}$. If we take the highest such $\bar{\rho}$ and $\bar{\lambda}$, then P_l never sends reports $(0, \rho > \bar{\rho})$ and P_r never sends reports $(\lambda > \bar{\lambda}, 0)$. By the same argument as in the proof of Proposition 3 we have that η^*_ρ is strictly increasing in ρ for $\rho > \bar{\rho}$, and η^*_λ is strictly decreasing in λ for $\lambda > \bar{\lambda}$. Moreover, by continuity of $\eta(\lambda, \rho)$ for $\rho > \bar{\rho}$ we have $\eta^*_\rho = \frac{\rho}{\rho+\lambda_\rho}$ and for $\lambda > \bar{\lambda}$ we have $\eta^*_\lambda = \frac{\rho_\lambda}{\rho_\lambda+\lambda}$. The definition of η^*_ρ is the same as in equation (3), η^*_λ is defined analogously and they are unique by the same argument as in Proposition **3** .

Step 6 There are three possible situations, $\eta^* > \frac{\bar{\rho}}{\bar{\rho}+\bar{\lambda}}$, $\eta^* < \frac{\bar{\rho}}{\bar{\rho}+\bar{\lambda}}$ or $\eta^* = \frac{\bar{\rho}}{\bar{\rho}+\bar{\lambda}}$. Below I describe the shape of the equilibrium if $\eta^* < \frac{\bar{\rho}}{\bar{\rho}+\bar{\lambda}}$. For the remaining cases the discussion is analogous.

If $\eta^* < \frac{\bar{\rho}}{\bar{\rho}+\hat{\lambda}}$, then for all (λ, ρ) that lie in the shaded area in the figure below, and only for these reports (or when $\frac{R}{L+R} = \eta^*$), we have $\eta(\lambda, \rho) = \eta^*$.

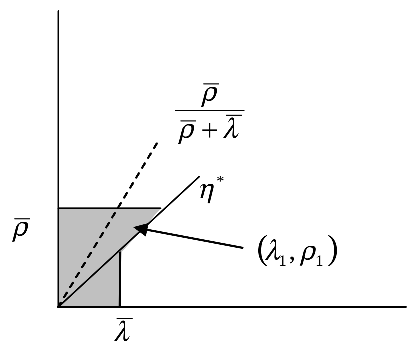

To see this notice that by Step 3 $\eta(\lambda, \rho) = \eta^*$ for all $(\lambda, \rho) : \lambda \leq \bar{\lambda}$ and $\rho \leq \bar{\rho}$. Take report $(\lambda_1, \rho_1) : \frac{\rho_1}{\rho_1+\lambda_1} > \eta^*$ like in the figure above. By Step 4, $\eta(\lambda_1, \rho_1) \leq \eta^*$. Assume that $\eta(\lambda_1, \rho_1) < \eta^*$; then only P_l may send (λ_1, ρ_1), which by Step 4 implies that $\eta(\lambda_1, \rho_1) = \eta^*_{\lambda_1}$, and by Step 5 η^*_λ is strictly decreasing in λ, which in turn implies that $(\lambda_1, \rho_1) \in \sigma_{P_l}(L, R)$ only if $L = \lambda_1$ and $R \geq \rho_1$. But this implies that $\eta(\lambda_1, \rho_1) \geq \frac{\rho_1}{\rho_1+\lambda_1} > \eta^*$, which is a contradiction.

Step 6 allows us to summarize the shape of any equilibrium. If $\eta^* < \frac{\bar{\rho}}{\bar{\rho}+\bar{\lambda}}$, then the equilibrium looks like in the figure below.

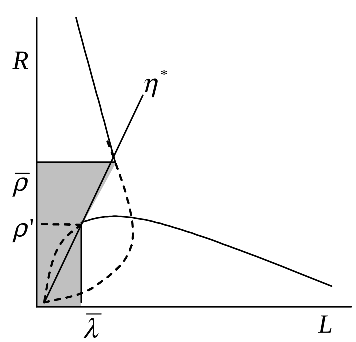

Figure A

The grey area represents all reports that generate the same belief η^*. The solid curves represent the areas in which only P_r and H (along the vertical axis) or only P_l and H (along the horizontal axis) send reports.

In what follows below I use $\bar{R}$ instead of $\bar{\rho}$ and $\bar{L}$ instead of $\bar{\lambda}$. In the proof of Proposition 3 I have shown that η^*_ρ (and therefore also η^*_λ) exists and is unique and strictly increasing for $\rho > \bar{R}$ (η^*_λ is strictly decreasing for $\lambda > \bar{L}$). I will show now that $\bar{R}$ and $\bar{L}$ and η^* are unique. Let all reports (λ, ρ) that generate η^* and either $\rho \leq \bar{R}$ or $\lambda \leq \bar{L}$ be called the double ambiguity area (DAA).

First, by continuity of $\eta(\lambda, \rho)$, $\bar{R}$ and $\bar{L}$ must be such that $\eta^*_{\bar{R}+\varepsilon} \rightarrow \eta^*$ and $\eta^*_{\bar{L}+\varepsilon} \rightarrow \eta^*$ as $\varepsilon \rightarrow 0$. This means that $\bar{R}$ and $\bar{L}$ must satisfy equations

$$\begin{aligned}\eta^*_\rho \equiv & \frac{\rho}{\rho+\lambda_\rho} = \Pr(H|\rho, \lambda \leq \lambda_\rho) E[q_R|R=\rho, L \leq \lambda_\rho] \\ & + \Pr(P_r|\rho, \lambda \leq \lambda_\rho) E[q_R|R=\rho],\end{aligned} \tag{7}$$

$$\begin{aligned}\eta^*_\lambda \equiv & \frac{\rho_\lambda}{\rho_\lambda+\lambda} = \Pr(H|\lambda, \rho \leq \rho_\lambda) E[q_R|L=\lambda, R \leq \rho_\lambda] \\ & + \Pr(P_l|\lambda, \rho \leq \rho_\lambda) E[q_R|L=\lambda],\end{aligned} \tag{8}$$

which can be rewritten as

$$\eta^* = \frac{\frac{\pi_r}{1-\ \pi_l}\int_0^{1-\bar{R}} \frac{\bar{R}}{\bar{R}+L} f^L\left(L|\bar{R}\right)dL + \frac{\pi_H}{1-\ \pi_l}\int_0^{\lambda_{\bar{R}}} \frac{\bar{R}}{\bar{R}+L} f^L\left(L|\bar{R}\right)dL}{\frac{\pi_r}{1-\ \pi_l} + \frac{\pi_H}{1-\ \pi_l} F\left(\lambda_{\bar{R}}|\bar{R}\right)}, \tag{9}$$

$$\eta^* = \frac{\frac{\pi_l}{1-\ \pi_r}\int_0^{1-\bar{L}} \frac{R}{R+\bar{L}} f^R\left(R|\bar{L}\right)dR + \frac{\pi_H}{1-\ \pi_r}\int_0^{\rho_{\bar{L}}} \frac{R}{R+\bar{L}} f^R\left(R|\bar{L}\right)dR}{\frac{\pi_l}{1-\ \pi_r} + \frac{\pi_H}{1-\ \pi_r} F^R\left(\rho_{\bar{L}}|\bar{L}\right)}. \tag{10}$$

These equations uniquely determine $\bar{R}$ and $\bar{L}$ as a function of η^*. For all reports in DAA to generate the same belief η^*, this belief must satisfy

$$\begin{aligned}\eta^* = & P(P_r|DAA) E\left[\frac{R}{R+L}|DAA, P_r\right] + \\ & + P(P_l|DAA) E\left[\frac{R}{R+L}|DAA, P_l\right] + P(H|DAA) E\left[\frac{R}{R+L}|DAA, H\right].\end{aligned}$$

Recall Figure A; the equation above can be rewritten as follows (where f is used instead of $f(L, R)$ to shorten the formula):

$$\begin{aligned}\eta^* = & \frac{\pi_r \int_0^{\bar{R}}\int_0^{1-R} q_R f dL dR + \pi_l \int_0^{\bar{L}}\int_0^{1-L} q_R f dL dR}{\pi_r \int_0^{\bar{R}}\int_0^{1-R} f dL dR + \ \pi_l \int_0^{\bar{L}}\int_0^{1-L} f dL dR + \pi_H\left(\int_{\frac{\eta^*\bar{L}}{1-\eta^*}}^{\bar{R}}\int_0^{\frac{(1-\eta^*)R}{\eta^*}} f dL dR + \int_0^{\frac{\eta^*\bar{L}}{1-\eta^*}}\int_0^{\bar{L}} f dL dR\right)} \\ & + \frac{\pi_H\left(\int_{\frac{\eta^*\bar{L}}{1-\eta^*}}^{\bar{R}}\int_0^{\frac{(1-\eta^*)R}{\eta^*}} q_R f dL dR + \int_0^{\frac{\eta^*\bar{L}}{1-\eta^*}}\int_0^{\bar{L}} q_R f dL dR\right)}{\pi_r \int_0^{\bar{R}}\int_0^{1-R} f dL dR + \ \pi_l \int_0^{\bar{L}}\int_0^{1-L} f dL dR + \pi_H\left(\int_{\frac{\eta^*\bar{L}}{1-\eta^*}}^{\bar{R}}\int_0^{\frac{(1-\eta^*)R}{\eta^*}} f dL dR + \int_0^{\frac{\eta^*\bar{L}}{1-\eta^*}}\int_0^{\bar{L}} f dL dR\right)}\end{aligned} \tag{11}$$

The LHS is continuous, strictly increasing and $LHS \in [0,1]$. The RHS is continuous, and as $\eta^* \rightarrow 0$, equations (9) and (10) imply that $\bar{R} \rightarrow 0$ and $\bar{L} \rightarrow 1$; therefore the

$RHS \to \frac{\int_0^1 \int_0^{1-L} \frac{R}{R+L} f dL dR}{\int_0^1 \int_0^{1-L} f dL dR} > 0$. Similarly, as $\eta^* \to 1$, by equations (9) and (10) $\bar{R} \to 1$ and $\bar{L} \to 0$; therefore, $RHS \to \frac{\int_0^1 \int_0^{1-R} \frac{R}{R+L} f dL dR}{\int_0^1 \int_0^{1-R} f dL dR} < 1$. Therefore, there exists η^* that solves equation (11). To show the uniqueness we can take the derivative of the RHS of equation (11) with respect to η^* and evaluate it at the point at which $\eta^* = RHS(\eta^*)$. We have $\frac{dRHS}{d\eta^*} = \frac{\partial RHS}{\partial \eta^*} + \frac{\partial RHS}{\partial \bar{R}} \frac{d\bar{R}}{d\eta^*} + \frac{\partial RHS}{\partial \bar{L}} \frac{d\bar{L}}{d\eta^*}$, and using equation (9) and equation (10) we can show that for $\eta^* = RHS(\eta^*)$ we have $\frac{\partial RHS}{\partial \eta^*} = 0$, $\frac{\partial RHS}{\partial \bar{R}} = 0$, and $\frac{\partial RHS}{\partial \bar{L}} = 0$. Every time $\eta^* = RHS(\eta^*)$, the derivative $\frac{dRHS}{d\eta^*} = 0$; therefore, there is at most one solution. ■

Proof of Proposition 5

Equation (6) characterizes η_R^*. It does not depend on π_l; therefore, we can take the limit of equation (6) keeping π_r constant. We get $\lim_{\pi_H \to 1} \eta_R^* = E\left[\frac{R}{R+L}|R\right]$, and $\lim_{\pi_H \to 0} \eta_R^* \to 1$. The definition of λ_R, $\eta_R^* \equiv \frac{R}{R+\lambda_R}$, implies that $\lim_{\pi_H \to 0} \lambda_R = 0$; that is, the reports of the persuader become more extreme, and that $\lim_{\pi_H \to 1} \lambda_R = \bar{\lambda}_R = \frac{1-E[q_R|R]}{E[q_R|R]} R$. The analogous holds for P_l.

If we keep $\frac{\pi_H}{1-\pi_r}$ constant, the shape of the ambiguity area for P_l remains unchanged, which can be seen if we investigate the analog of equation (6) for P_l. Keeping $\frac{\pi_H}{1-\pi_r}$ constant implies that π_H increases as π_r decreases; therefore, equation (6) implies that the ambiguity area of P_r shrinks. That means that $\lim_{\pi_r \to 0} \lambda_R \to 0$, which means that the reports of P_r become more extreme. Now, I will show that η^* increases as π_r goes down.

η^* is determined by equation (11), where $\bar{R}$ and $\bar{L}$ are determined by equations (9) and (10). If we take the derivative of the RHS of (11) with respect to π_r and evaluate at η^*, we get

$$\begin{aligned} \frac{dRHS}{d\pi_r}|_{\eta^*} &\overset{sign}{=} \frac{1}{(1-\pi_r)} \int_0^{\bar{R}} \int_0^{\infty} \left(\frac{R}{R+L} - \eta^*\right) f dL dR + \\ &+ \left(\pi_r \int_0^{\infty} \left(\frac{\bar{R}}{\bar{R}+L} - \eta^*\right) f dL + \pi_H \int_0^{\frac{1-\eta^*}{\eta^*}\bar{R}} \left(\frac{\bar{R}}{\bar{R}+L} - \eta^*\right) f dL\right) \frac{d\bar{R}}{d\pi_r} \\ &= \frac{1}{(1-\pi_r)} \int_0^{\bar{R}} \int_0^{\infty} \left(\frac{R}{R+L} - \eta^*\right) f dL dR < 0, \end{aligned}$$

where the last equality comes from using equation (9). Recall that η^* is the point of intersection of the $LHS\,(\eta)$ and the $RHS\,(\eta)$ of equation (11) for the initial π_r, like in the picture below.

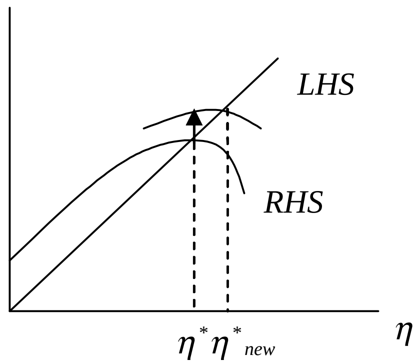

The fact that $\frac{dRHS}{d\pi_r} < 0$ implies that as π_r decreases the function $RHS(\eta)$ shifts up, which implies that the new $\eta^*_{new} > \eta^*$.

This implies that when π_r goes down, P_r is better-off, and P_l is worse-off. As π_r decreases, η^* and η^*_R for each R increase, which implies that for each R the persuader toward *Right* can induce a higher belief. Since the shape of the ambiguity area for P_l is not affected for big L, P_l can induce the same belief. However, the belief in the double ambiguity area is higher, and it is achieved for lower $\bar{L}_{new} < \bar{L}$ which means that for $L < \bar{L}$, P_l induces higher beliefs than before.

Showing that the utility of the decision maker increases requires some tedious algebra, which I omit here, but the result is intuitive, since it is more likely that the decision maker faces the honest expert. ■

Proof of Proposition 6

Step 1 For a given q_R, the expected utility of the decision maker with parameter θ is

$$\text{for } q_R < \theta : U(q_R) = (1-\pi)\int_0^{\infty}(\theta - q_R)\, g(N;z)\, dN + \pi \int_0^{\frac{R_\theta}{q_R}}(\theta - q_R)\, g(N;z)\, dN$$
$$+\pi \int_{\frac{R_\theta}{q_R}}^{\infty}(q_R - \theta)\, g(N;z)\, dN = (1-2\pi)(\theta - q_R) + 2\pi(\theta - q_R)\, G\left(\frac{R_\theta}{q_R}; z\right),$$

$$\begin{aligned}\text{for } q_R &> \theta : U(q_R) = \int_0^{\frac{R_\theta}{q_R}}(\theta - q_R)\, g(N;z)\, dN + \int_{\frac{R_\theta}{q_R}}^{\infty}(q_R - \theta)\, g(N;z)\, dN = \\ &= 2(\theta - q_R)\, G\left(\frac{R_\theta}{q_R}; z\right) + (q_R - \theta).\end{aligned}$$

Hence

$$
\begin{aligned}
E[U] &= \int_0^\theta \left((1-2\pi)(\theta - q_R) + 2\pi(\theta - q_R) G\left(\frac{R_\theta}{q_R}; z\right)\right) dq_R \\
&\quad + \int_\theta^1 \left(2(\theta - q_R) G\left(\frac{R_\theta}{q_R}; z\right) + (q_R - \theta)\right) dq_R.
\end{aligned}
$$

Step 2 R_θ is such that the decision maker is indifferent between *Right* and *Left;* therefore, $\frac{dE[U]}{dz} = \frac{\partial E[U]}{\partial z} + \frac{\partial E[U]}{\partial R_\theta}\frac{dR_\theta}{dz} = \frac{\partial E[U]}{\partial z}$. Hence, we can look only at the direct effect of changing z.

We have

$$
\frac{\partial E[U]}{\partial z} = \int_0^\theta 2\pi(\theta - q_R) G_z\left(\frac{R_\theta}{q_R}; z\right) dq_R + \int_\theta^1 2(\theta - q_R) G_z\left(\frac{R_\theta}{q_R}; z\right) dq_R.
$$

I will prove in Step3 that $\bar{N} > \frac{R_\theta}{\theta}$. Using this, the second expression is positive since for $q_R \in (\theta, 1)$ we have $(\theta - q_R) < 0$ and $\frac{R_\theta}{q_R} \leq \frac{R_\theta}{\theta} < \bar{N}$. The first expression can be rewritten

$$
\begin{aligned}
&\int_0^\theta 2\pi(\theta - q_R) G_z\left(\frac{R_\theta}{q_R}; z\right) dq_R = \\
= &\int_0^{\frac{R_\theta}{\bar{N}}} 2\pi(\theta - q_R) G_z\left(\frac{R_\theta}{q_R}; z\right) dq_R + \int_{\frac{R_\theta}{\bar{N}}}^\theta 2\pi(\theta - q_R) G_z\left(\frac{R_\theta}{q_R}; z\right) dq_R,
\end{aligned}
$$

where the fist expression is positive and the second is negative.

The fact that the mean of N is unaffected by changes in z implies that $\int_0^\infty G_z(N; z) = 0$, which in turn implies that there exists a set (a, b) with $a \geq \bar{N}$ such that

$$
\int_{\frac{R_\theta}{\bar{N}}}^\theta \left(-G_z\left(\frac{R_\theta}{q_R}; z\right)\right) dq_R = \int_{\frac{R_\theta}{y}}^{\frac{R_\theta}{x}} G_z\left(\frac{R_\theta}{q_R}; z\right) dq_R.
$$

Clearly $\frac{R_\theta}{x} < \frac{R_\theta}{\bar{N}}$, hence to complete the proof we need to show that

$$
\int_{\frac{R_\theta}{y}}^{\frac{R_\theta}{x}} 2\pi(\theta - q_R) G_z\left(\frac{R_\theta}{q_R}; z\right) dq_R > \int_{\frac{R_\theta}{\bar{N}}}^\theta 2\pi(\theta - q_R)\left(-G_z\left(\frac{R_\theta}{q_R}; z\right)\right) dq_R.
$$

But

$$\int_{\frac{R_\theta}{y}}^{\frac{R_\theta}{x}} 2\pi\left(\theta - q_R\right) G_z\left(\frac{R_\theta}{q_R};z\right) dq_R > 2\pi\left(\theta - \frac{R_\theta}{x}\right)\int_{\frac{R_\theta}{y}}^{\frac{R_\theta}{x}} G_z\left(\frac{R_\theta}{q_R};z\right) dq_R >$$

$$> 2\pi\left(\theta - \frac{R_\theta}{\bar N}\right)\int_{\frac{R_\theta}{y}}^{\frac{R_\theta}{x}} G_z\left(\frac{R_\theta}{q_R};z\right) dq_R = 2\left(\theta - \frac{R_\theta}{\bar N}\right)\int_{\frac{R_\theta}{\bar N}}^{\theta}\left(-G_z\left(\frac{R_\theta}{q_R};z\right)\right) dq_R >$$

$$> \int_{\frac{R_\theta}{\bar N}}^{\theta} 2\left(\theta - q_R\right)\left(-G_z\left(\frac{R_\theta}{q_R};z\right)\right) dq_R.$$

Step 3 $\bar N > \frac{R_\theta}{\theta}$

Given the assumptions, the conditional p.d.f. of N given R is $g\left(N|R;z\right) = \frac{\frac{g(N;z)}{N}}{\int_R^1 \frac{1}{N} g(N;z)dN}$. The number of arguments in favor of *Right* for which the decision maker with parameter θ is indifferent between the alternatives, R_θ, is defined by $\theta = \eta^*_{R_\theta}$, and by the definition of $\eta^*_{R_\theta}$, one obtains

$$\begin{aligned} \theta &= \eta^*_{R_\theta} \geq E\left[\frac{R_\theta}{N}\right] = \frac{\int_{R_\theta}^1 \frac{R_\theta}{N}\frac{1}{N} g\left(N;z\right)dN}{\int_{R_\theta}^1 \frac{1}{N} g\left(N;z\right)dN} = \\ &= \frac{R_\theta E\left[\frac{1}{N^2}|N > R_\theta;z\right]}{E\left[\frac{1}{N}|N > R_\theta;z\right]} > R_\theta E\left[\frac{1}{N}|N > R_\theta;z\right] > \frac{R_\theta}{\bar N} \end{aligned}$$

Full revelation of information in the case of $z = \infty$ is a straightforward result. ■

Proof of Proposition 7

This proof shows that for any fixed R, all types of the decision maker are at least as well-off in expectation in the continuous equilibria as in any discontinuous one. Let $\eta^{*c}_R \equiv \eta\left(0,R\right)$ in any continuous equilibrium and let $\eta^*_R \equiv \eta\left(0,R\right)$ in some discontinuous equilibrium. Recall that λ_R is the highest λ such that $\eta\left(\lambda,R\right) = \eta^*_R$ for all $\lambda < \lambda_R$, and let λ^c_R be the highest λ such that $\eta\left(\lambda,R\right) = \eta^{*c}_R$ for all $\lambda < \lambda^*_R$. From the proof of Propositions 2 and 3 we know that $\frac{R}{R+\lambda^c_R} = \eta^{*c}_R$ and $\frac{R}{R+\lambda_R} \leq \eta^*_R$. Clearly, for a given R and keeping η^*_R fixed, the decision maker is weakly better-off when $\frac{R}{R+\lambda_R} = \eta^*_R$ than when $\frac{R}{R+\lambda_R} < \eta^*_R$, as in both cases she makes the same decision when facing P_r, but in the latter case she may make worse decisions when facing H. Hence I assume that $\frac{R}{R+\lambda_R} = \eta^*_R$, and show that even in this best scenario for the discontinuous equilibria the decision maker still prefers the continuous ones.

First, recall that η_R^{*c} is defined as x that solves the following equation:

$$x = \frac{\pi E\left[q_R|R\right] + (1-\pi)\Pr\left(q_R > x\right) E\left[q_R|q_R > x, R\right]}{\pi + (1-\pi)\Pr\left(q_R > x|R\right)}.$$

At the end of the proof of Proposition 3 I have shown that $RHS < LHS$ for $x < \eta_R^{*c}$ and $RHS > LHS$ for $x > \eta_R^{*c}$; in other words

$$\frac{\pi E\left[q_R|R\right] + (1-\pi)\Pr\left(q_R > x\right) E\left[q_R|q_R > x, R\right]}{\pi + (1-\pi)\Pr\left(q_R > x|R\right)} \begin{cases} > x \text{ if } x < \eta_R^{*c} \\ < x \text{ if } x > \eta_R^{*c} \end{cases}. \tag{12}$$

For each R, we have three cases: 1. $\eta_R^* < \eta_R^{*c}$, 2. $\eta_R^* > \eta_R^{*c}$ or 3. $\eta_R^* = \eta_R^{*c}$.

Case 1: $\eta_R^* < \eta_R^{*c}$. A decision maker with $\theta \geq \eta_R^{*c}$ or with $\theta \leq \eta_R^*$ makes the same decision in both equilibria. For a given R a decision maker with $\theta \in \left(\eta_R^*, \eta_R^{*c}\right)$ chooses *Left* in the discontinuous equilibrium for all L, but *Right* in the continuous equilibria when the expert is a persuader for all L, or when the expert is honest and $L > \theta$. The expected quality q_R conditional on the events in which she chooses *Right* in the continuous equilibrium is

$$\frac{\pi E\left[q_R|R\right] + (1-\pi) P\left(q_R > \theta\right) E\left[q_R|q_R > \theta, R\right]}{\pi + (1-\pi) P\left(q_R > \theta|R\right)},$$

which by equation 12 is greater than θ, hence *Right* is the better choice. This means that the decision maker is better off in the continuous equilibrium for this R.

Case 2: $\eta_R^* > \eta_R^{*c}$. A decision maker with $\theta \leq \eta_R^{c*}$ or with $\theta \geq \eta_R^*$ makes the same decision in both equilibria. For a given R a decision maker with $\theta \in \left(\eta_\rho^{*c}, \eta_\rho^*\right)$ chooses *Right* in the discontinuous equilibrium for all L, but *Left* in the continuous equilibria when the expert is a persuader for all L, or when the expert is honest and $L > \theta$. The expected quality q_R conditional on the events in which she chooses *Left* in the continuous equilibrium is

$$\frac{\pi E\left[q_R|R\right] + (1-\pi) P\left(q_R > \theta\right) E\left[q_R|q_R > \theta, R\right]}{\pi + (1-\pi) P\left(q_R > \theta|R\right)},$$

which by equation 12 is smaller than θ, hence *Left* is the better choice. This means the decision maker is again better off in the continuous equilibrium for this R.

Case 3: $\eta_R^* = \eta_R^{*c}$. It is immediate that the decision maker is indifferent between the equilibria. ■

References

[1] Anderson, Thomas W. and Linda L. Golden (1984), "Bank Promotion Strategy," *Journal of Advertising Research,* 24, 53-65.

[2] Battaglini, Marco (2002), "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica,* 70, 1379-1401.

[3] Chen, Ying (2007), "Perturbed Communication Games with Honest Senders and Naive Receivers," SSRN Working Paper.

[4] Crawford, Vincent P. and Joel Sobel (1982), "Strategic Information Transmission," *Econometrica,* 50, 1431-1451.

[5] Crowley, Ayn E. and Wayne D. Hoyer (1994), "An Integrative Framework for Understanding Two-sided Persuasion," *Journal of Consumer Research,* 20, 561-574.

[6] Darley, John M. and Paget H. Gross (1983), "A Hypothesis-Confirming Bias in Labeling Effects," *Journal of Personality and Social Psychology,* 44, 20-33.

[7] Dye, Ronald A. (1985), "Disclosure of Nonproprietary Information," *Journal of Accounting Research,* 23, 123-145.

[8] Fishman, Michael and Kathleen Hagerty (2003), "Mandatory versus Voluntary Disclosure in Markets with Informed and Uninformed Customers," *Journal of Law, Economics, and Organization,* 19, 45-63.

[9] Glazer, Jacob and Ariel Rubinstein (2001), "Debates and Decisions: On a Rationale of Argumentation Rules," *Games and Economic Behavior,* 36, 158-173.

[10] Glazer, Jacob and Ariel Rubinstein (2004), "On Optimal Rules of Persuasion," *Econometrica,* 72, 1715-1736.

[11] Grossman, Sanford J. (1981) "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics,* 24, 561-483.

[12] Grubb, Michael D. (2007), "Developing a Reputation for Reticence," SSRN Working Paper.

[13] Hovland, Carl I., Arthur A. Lumsdaine, and Frederick D. Sheffield, *Experiments on Mass Communication*, Princeton University Press, 1949, 201-227.

[14] Jennings, Dennis L., Mark R. Lepper, and Lee Ross (1981), “Persistence of Impressions of Personal Persuasiveness: Perseverance of Erroneous Self-Assessments outside the Debriefing Paradigm,” *Personality and Social Psychology Bulletin,* 7, 257–263.

[15] Jovanovic, Boyan (1982), "Truthful Disclosure of Information," *The Bell Journal of Economics,* 13, 36-44.

[16] Kartik, Navin (2007), "Strategic Communication with Lying Costs," http://www.columbia.edu/~nk2339/Papers/sclc.pdf.

[17] Kartik, Navin, Marco Ottaviani and Francesco Squintani (2007), "Credulity, Lies, and Costly Talk," *Journal of Economic Theory,* 134, 93-116.

[18] Krishna, Vijay and John Morgan (2001), "A Model of Expertise," *The Quarterly Journal of Economics,* 116, 747-775.

[19] Lord, Charles. G, Lee Ross and Mark R. Lepper (1979), "Biased Assimilation adn Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence," *Journal of Personality and Social Psychology, 37, 2098-2109.*

[20] Lumsdaine, Arthur A. and Irving L. Janis (1953), "Resistance to ”Counterpropaganda’ Produced by One-sided and Two-sided ”Propaganda’ Presentations," *The Public Opinion Quarterly,* 17, 311-318.

[21] Matthews, Steven A. and Andrew Postlewaite (1985), "The Economics of Quality Testing and Disclosure," *The RAND Journal of Economics,* 16, 328-340.

[22] Milgrom, Paul R. (1981), "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics,* 12, 380-391.

[23] Milgrom, Paul R. and John Roberts (1986), "Relying on the Information of Interested Parties," *The RAND Journal of Economics,* 17, 18-32.

[24] Pechmann, Cornelia (1992), "Predicting When Two-Sided Ads Will Be More Effective Than One-Sided Ads: The Role of Correlational and Correspondent Inferences," *Journal of Marketing Research,* 29, 441-453.

[25] Plous, Scott (1991), “Biases in the Assimilation of Technological Breakdowns: Do Accidents Make Us Safer?” J*ournal of Applied Social Psychology,* 21, 1058-1082.

[26] Rabin, Matthew and Joel L. Schrag (1999), "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics,* 114, 37-82.

[27] Shin, Hyun Song (1994a), "The Burden of Proof in a Game of Persuasion," *Journal of Economic Theory,* 64, 253-264.

[28] Shin, Hyun Song (1994b), "News Management and the Value of Firms," *The RAND Journal of Economics,* 25, 58-71.

[29] Shin, Hyun Song (2003), "Disclosures and Asset Returns," *Econometrica,* 71, 105-133.

[30] Smith, Robert E. and Shelby D. Hunt (1978), "Attributional Processes and Effects in Promotional Situations," *Journal of Consumer Research,* 5, 149-158.

[31] Verrecchia, Robert E. (2001), "Essays on disclosure," *Journal of Accounting and Economics,* 32, 97–180.

[32] Viscusi, Kip W. (1978), "A Note on ”Lemons’ Markets with Quality Certification," *The Bell Journal of Economics,* 9, 277-279.

[33] Wolinsky, Asher (2003), "Information Transmission When the Sender’s Preferences are Uncertain," *Games and Economic Behavior,* 42, 319-326.