Experts, Conflicts of Interest, and Reputation for Ability*

Filippo Pavesi[a,†], Massimo Scotti[b,‡]

[a]Department of Economics, University of Milan-Bicocca.

[b]Economics Discipline Group, University of Technology, Sydney.

December 2011

Abstract

We analyze a model of cheap talk in which an expert that faces a conflict of interest with a decision maker is concerned about establishing a reputation for having accurate information. In this environment, the incentive of the expert to establish a reputation for competence has a non-monotonic effect on the degree of information revelation. An increase in reputation above a certain threshold always makes truthful revelation more difficult to achieve. This is driven by the fact that experts with greater reputation for ability can more easily sway the beliefs of decision makers in a desired direction. Thus,

*We would like to thank Pierpaolo Battigalli, Paolo Colla, Amil Dasgupta, Giovanna Iannantuoni, Andrea Pratt, Jean-Charles Rochet, and Guido Tabellini. We also thank Stefano Bonini and Abigail Brown, and all participants at the NASMES 2011 in St. Louis, the EEA 2010 Congress in Glasgow, the Asset 2010 Conference in Alicante, and the SERC 2009 in Singapore for helpful comments. Financial support from Universita' L. Bocconi, the School of Finance and Economics at UTS Sydney, the Department of Economics at University of Milano-Bicocca, and the Department of Public Policy at CEU Budapest is gratefully acknowledged. All errors remain our own.

†Email: filippo.pavesi@unimib.it

‡Corresponding author at: Economics Discipline Group, University of Technology, Sydney, PO Box 123 Broadway NSW 2007 Australia. Phone number: +61 2 95147776. Email: massimo.scotti@uts.edu.au

higher levels of reputation exacerbate the incentives of biased experts to misreport their private information. Decision makers will therefore be better off consulting less reputable experts when conflicts are more pronounced.

Keywords: Experts; Reputation; Cheap Talk; Conflicts of Interest; Information Transmission.

JEL Classification: C72, D82, D83

1 Introduction

There are several economic and political settings in which an expert that is called on to provide information to a decision maker faces an intrinsic conflict of interest. In many such cases, it is plausible to presume that the decision maker is aware of this conflict, at least to some extent. For example, many investors are likely to have a good understanding of the fact that financial analysts have incentives to provide biased reports.[1] Similarly, in the political arena, the electoral body is likely to know that government agencies have reasons to bias their macroeconomic forecasts towards those that favor politicians.[2] A standard argument is that the concern of an expert about establishing a reputation for being competent would mitigate this conflict.[3]

The central question that we address in this paper is how reputation for ability and more in general the perceived quality of an expert's information affect the communication process in a context in which the expert's bias is commonly known. We find that in the presence of

[1] See for example Michaeli and Womack (1999) and Barber et al. (2006, 2007) showing that affiliated analysts have an optimism bias resulting from their involvement in the investment banking activity of their brokerage house.

[2] Weatherford (1987), Alesina and Roubini (1997) and Carlsen (1999) document that incumbent governments generally prefer agencies that are more inclined to provide optimistic forecasts. In these cases, the conflict of interest originates from the ability of the executive branch to sanction agencies that fail to act in its interest by proposing budget cuts, disposing of executives or even advocating termination of the agency.

[3] See for example Mikhail et al. (1999), Hong and Kubik (2003) and Fang and Yasuda (2009) for the case of financial analysts. Heclo (1975), Rourke (1992), Carpenter (2001), Wilson (1989), Bendor et al. (1985) and Banks and Weingast (1992) document the disciplining role of reputation and career concerns in the political arena.

conflicts of interest, decision makers are not necessarily better off when they consult more reputable experts (i.e., experts with a better expected quality of information). Indeed, we show that it may be optimal for a decision maker to consult a less reputable expert precisely when incentives are less aligned.

We derive these conclusions in a model of cheap talk in which an expert privately observes a (binary) signal about a (binary) state of the world, and subsequently makes a cheap talk report to a decision maker. The accuracy of the signal depends on the ability of the expert, which is unknown both to the expert and the decision maker. The decision maker observes the report of the expert and updates his belief about the state. Once the state has been publicly revealed, the decision maker uses the report to also update his belief about the ability of the expert. To capture the presence of conflicts of interest and the expert's reputational concern for ability, we assume that the payoff of the expert is increasing both in the decision maker's belief that the state is high and in the decision maker's belief about the ability of the expert. Thus, when deciding which report to make, the expert trades off the reputational reward of providing a correct report against the benefit of using his credibility to sway the receiver's beliefs about the state in the desired direction. The payoff function of the expert is assumed to be common knowledge. This implies that the decision maker is aware of the expert's bias.

Our first result is that an increase in the level of initial reputation has a non-monotonic impact on the expert's incentives to truthfully reveal his information. In particular, beyond a certain threshold, any increase in initial reputation always reduces informational efficiency. This is in contrast with the case of no conflicts of interest, where an increase in the level of initial reputation always has a positive effect on the amount of information revealed. Intuitively, as reputation becomes sufficiently high, there is less scope for reputation acquisition, and reputation becomes less effective in disciplining the bias of the expert.[4] At the same time,

[4] This effect is consistent with the reputational incentives identified by Holmström (1999) in a model of effort provision, where managers exert more effort in the initial stages of their career, when uncertainty on

as the reputation of an expert increases, the decision maker gives more weight to the advice of the expert. This in turn provides a biased expert with a higher incentive to misreport. Thus, a distinctive feature of a model with conflicts of interest is that the expert's incentive to misreport increases endogenously with the level of the expert's reputation. This has two implications. First, ceteris paribus, the distortionary incentives provided by reputation may induce a more reputable expert to lie when a less reputable expert does not, suggesting that in the presence of conflicts of interest decision makers are not necessarily better off when they consult more reputable experts. Second, as conflicts of interest become more severe, informational efficiency is maximized at progressively lower values of reputation, implying that it may be optimal for decision makers to consult less reputable experts precisely when incentives are less aligned.

A second finding is that an increase in the accuracy of the expert's signals may lead to a reduction in informational efficiency in spite of its overall positive effect on the expected quality of the expert's information. This is the case when the improvement in the expected quality of the expert's information comes entirely from an increase in the accuracy of the signal of the less talented expert. In this case, as the abilities of the experts converge, the reputational gain of being recognized as a good expert tends to fade thereby reducing the disciplining role of reputation. At the same time, the improved quality of information enhances the credibility of advice, increasing the returns from biased reports. This result never arises in a setting without conflicts, where an improvement in the accuracy of expert's signals unambiguously increases informational efficiency.

Finally, we show that in the presence of conflicts of interest and reputational concerns for ability, truthful revelation becomes possible only when public information is rather contrary to the state towards which the expert wishes to sway the belief of the decision maker. For example, in our binary model, an expert with a strong bias towards the high state will report truthfully only when the prior probability of the high state is relatively small. This

their ability is higher and the scope for reputation acquisition is greater.

result, which again arises from the interaction between the expert's conflict of interest and his reputational concerns, suggests that a biased expert should be consulted over issues for which the public consensus is polarized around a belief that is opposite to the one that the expert would like to induce.

Our paper is related to two main strands of the literature on sender-receiver models of information transmission. The first strand analyzes information transmission in the case in which senders' and receivers' preferences are misaligned (Crawford and Sobel 1982, Sobel 1985; Benabou and Laroque 1992; Morris 2001). The second deals with experts that do not have explicit conflicts of interest with decision makers and are exclusively concerned about establishing a reputation for having accurate information (Ottaviani and Sorensen 2001, 2006; Trueman 1994).

A standard result of the first strand is that only noisy information can be credibly transmitted if the expert and the decision maker have conflicting preferences. In particular, the more biased the expert is, the noisier the information revealed (Crawford and Sobel 1982). Starting with Sobel (1985), this literature has analyzed games of cheap-talk in which there is uncertainty on the preferences of the expert, and the expert can establish a reputation for being unbiased (Benabou and Laroque 1992; Morris 2001). In particular, Morris (2001) highlights a potentially distortionary effect of reputation by showing that an advisor with preferences aligned with those of the decision maker may in fact distort his private information in order to build a reputation for being unbiased.[5] In our model, the preferences of the expert are assumed to be common knowledge and uncertainty is about the forecasting ability of the expert. The issue we address is whether experts with a higher reputation for competence are more likely to credibly transmit their information when it is well known that they are biased.

Our paper is closely related to Ottaviani and Sorensen (2001, 2006). They study infor-

[5] Ely and Valimaki (2003) obtain a result in the spirit of Morris (2001) in a principal-agent model whitout cheap-talk.

mation reporting by privately informed experts who are solely motivated by the desire to be perceived as competent, and show that honesty is impossible under very general conditions. In their model, the amount of information that is credibly transmitted is always increasing in the quality of the expert's information. By introducing conflicts of interest in a setting with reputation for ability, we show that greater quality of information is not necessarily associated with less misreporting. In some respect, our work is complementary to Bourjade and Jullien (2011) who also consider the case of a biased expert with a reputational concern for competence, but in a setting in which the expert has hard information. They consider strategic concealment of private information while we analyze the issue of misrepresentation of information.

The remainder of the paper is organized as follows. In Section 2, we introduce the general setup of the model. Section 3 characterizes the most informative equilibrium and analyzes the conditions under which truthtelling is possible, highlighting the incentives that lead experts to deviate from truthtelling. Section 4 examines how informational efficiency is affected by variations in the quality of the expert's information. Section 5 studies how changes in the intensity of the expert's bias affect the expert's incentives to report truthfully. Section 6 concludes.

2 The Model

An expert is called upon to provide information to a decision maker who has to make a forecast about the state of world. The state of the world ω is either high or low, i.e., $\omega \in \{h, l\}$, and all players hold the same prior belief θ that the state is h. At the beginning of the game, the expert observes a private and non-verifiable signal $s_i \in \{s_h, s_l\}$ about the true state, whose accuracy depends on the expert's ability t. We assume that the expert is either good or bad, i.e., $t \in \{g, b\}$, and that ability affects the accuracy of the signal as follows:

$$\Pr(s_h|t = g, \omega = h) = \Pr(s_l|t = g, \omega = l) = p,\ p \in (1/2, 1)\,, \tag{1}$$

$$\Pr(s_h|t = b, \omega = h) = \Pr(s_l|t = b, \omega = l) = z,\ z \in (1/2, p]\,. \tag{2}$$

Therefore, both expert types can count on an informative (yet imperfect) signal, with the good type having a more accurate signal than a bad type.[6] We assume that neither the expert nor the decision maker know the expert's type, and all players hold the same prior belief α that the expert is good.[7] We interpret α as the prior reputation for ability of the expert.

After observing the signal, the expert chooses to release a report to the decision maker in the form of a costless binary message $m_j \in \{m_h, m_l\}$. The decision maker observes message m_j and revises his beliefs about the true state of the world. We denote with $\widehat{\theta}_{\alpha,m_j} \equiv \Pr(\omega = h|m_j)$, the decision maker's posterior belief that the state of the world is h, given that message m_j was sent by the expert with prior reputation α. As we will see, in an equilibrium where some information is transmitted, the higher the reputation of the expert, the more the decision maker will trust the message sent. The subscript α highlights this relationship.

At the end of the game, the true state of the world is revealed and together with the message of the expert is used by the decision maker to revise his beliefs about the expert's ability.[8] We denote with $\widehat{\alpha}_{\omega,m_j} \equiv \Pr(t = g|\omega, m_j)$, the decision maker's posterior belief that

[6] All the results also hold for $z = \frac{1}{2}$.We make use of informative signals of bad types of experts, $z \in (1/2, p]$ in section (4.2) when we analyze how variations in z affect the amount of information that is transmitted in equilibrium.

[7] This assumption is without loss of generality as far as the key results of paper are concerned, and makes the analysis more tractable. When the expert knows his own type, in the most informative equilibrium both types of experts truthfully report their signal for intermediate values of the prior on the state. Outside the truthtelling region, partial revealing equilibria exist in which the good expert always reports truthfully and the bad expert lies with positive probability. Despite the different nature of the most informative equilibrium, all our results extend to the case of known ability.

[8] In fact, in our model the receivers perform the task of forecasting the state of the world and the expert's

the expert is good upon observing state ω and message m_j. We interpret $\widehat{\alpha}_{\omega,m_j}$ as the new level of reputation for ability acquired by the expert at the end of the game.

To model the expert's concern about establishing a reputation for being a valuable provider of information and the contemporaneous existence of conflicts of interest, we construct a game where the payoff of the expert depends positively on the decision maker's posterior beliefs $\widehat{\theta}_{\alpha,m_j}$ and $\widehat{\alpha}_{\omega,m_j}$, as follows:

$$\pi(m_j) = k\widehat{\theta}_{\alpha,m_j} + (1-k)\widehat{\alpha}_{\omega,m_j} \text{ with } k \in [0,1]\,. \tag{3}$$

The component $\widehat{\alpha}_{\omega,m_j}$ captures the concern of the expert to be perceived as having accurate information.[9] The component $\widehat{\theta}_{\alpha,m_j}$ gives the expert an incentive to inflate the decision maker's belief that the state is h, and thus creates a conflict of interest with the decision maker, since the expert now has a bias in favor of information that increases the decision maker's perception that the state is h.[10] Finally, the parameter $k \in [0,1]$ weighs these two components and can be seen as a measure of the severity of conflicts of interest. The structure and the parameters of the game (with the sole exception of the expert's signal) are common knowledge.[11]

Notice that interpreting h and l respectively as favorable and unfavorable states for the decision maker, the model represents the over-optimism bias that has been discussed both in the finance literature on sell side analysts and in the political science literature on government

ability. Notice that we do not explicitly model the payoff of the receivers. Instead, we follow the approach of Ottaviani and Sorensen (2006) and implicitly assume that receivers are rewarded for accurately forcasting both the state of the world and the ability of the expert.

[9] This reduced form to account for reputational concerns is widely adopted in studies that model the reputation of experts and managers (see for example Sharfstein and Stein (1990), Ottaviani and Sorensen (2006) and Gentzkow and Shapiro (2006)).

[10] Formally this game falls in the class of psychological games since the sender's payoff depends on the receiver's belief (Battigalli and Dufwenberg 2009).

[11] It is worth noticing that since also k is common knowledge, we do not address the case when receivers are uncertain about the incentives of the expert. See Sobel (1985), Benabou and Laroque (1992), and Morgan and Stocken (2003) for a formal analysis of the case in which there is uncertainty about the expert's preferences.

agencies' forecasts.[12] For the sake of exposition, in the remainder of the paper we will adopt this interpretation and refer to the expert's bias as to the over-optimism bias.

3 Equilibrium Analysis

In this section, we analyze the incentives of an expert to truthfully report his information and characterize the most informative equilibrium.[13]

At the moment of sending message m_j, the true state of the world is unknown to the expert. The expert uses his signal s_i to compute the expected impact of message m_j on his reputation, as follows:

$$E\left(\widehat{\alpha}_{\omega,m_j}|s_i\right) = \Pr(\omega = h|s_i)\widehat{\alpha}_{h,m_j} + \Pr(\omega = l|s_i)\widehat{\alpha}_{l,m_j}.$$

Therefore, the expected payoff of the expert from sending message m_j reads:

$$E\left(\pi(m_j)|s_i\right) = k\widehat{\theta}_{\alpha,m_j} + (1-k)E\left(\widehat{\alpha}_{\omega,m_j}|s_i\right).$$

Before analyzing the incentives of an expert to truthfully report his information, it is convenient to gain an intuition of the tensions involved in the reporting decision of the expert. In any equilibrium where some information is transmitted we have that $\widehat{\theta}_{\alpha,m_h} > \widehat{\theta}_{\alpha,m_l}$.[14] This introduces an incentive to report message m_h and represents a threat to truthtelling whenever signal s_l is received. In fact, the presence of reputational concerns counterbalances this over-optimism bias. As long as $k \in (0,1)$, the expert has to trade off the temptation of

[12] Assuming that the expert has an interest in inflating the receivers' belief about the state being h, is without loss of generality. Our setup is well suited for analyzing a more general setting, where the expert has an incentive to manipulate the receivers' beliefs in a desired direction.

[13] Our model presents the well-known problem of equilibrium multiplicity that is common to any cheap-talk game. A babbling equilibrium where all messages are taken to be meaningless and ignored always exists.

[14] Since the expert's signals are informative, in any equilibrium where signals are truthfully reported with some positive probability, the messages of the expert contain some information.

sending m_h with the negative effects that this message might have on his reputation in case the message turns out to be incorrect.

The equilibrium concept we use is that of Perfect Bayesian Equilibrium (PBE). The expert will truthtfully report signal s_i if and only if the expected payoff of truthtelling is greater than the payoff of reporting a message that is different from the signal received. Thus, a truthtelling equilibrium exists if and only if for every $i, j \in \{h, l\}$ with $i \neq j$, $E(\pi(m_i)|s_i) \geq E(\pi(m_j)|s_i)$, or equivalently:

$$k\widehat{\theta}_{\alpha,m_l} + (1-k)E(\widehat{\alpha}_{\omega,m_l}|s_l) \geq k\widehat{\theta}_{\alpha,m_h} + (1-k)E(\widehat{\alpha}_{\omega,m_h}|s_l), \tag{4}$$

$$k\widehat{\theta}_{\alpha,m_h} + (1-k)E(\widehat{\alpha}_{\omega,m_h}|s_h) \geq k\widehat{\theta}_{\alpha,m_l} + (1-k)E(\widehat{\alpha}_{\omega,m_l}|s_h). \tag{5}$$

In a truthtelling equilibrium, posterior reputation takes on only two possible values, which we denote with $\underline{\alpha}$ and $\overline{\alpha}$, where:

$$\underline{\alpha} \equiv \widehat{\alpha}_{l,m_h} = \widehat{\alpha}_{h,m_l},$$

$$\overline{\alpha} \equiv \widehat{\alpha}_{h,m_h} = \widehat{\alpha}_{l,m_l},$$

with $\overline{\alpha} > \alpha > \underline{\alpha}$.[15] Making a correct evaluation increases the expert's reputation from its initial level α to the higher level $\overline{\alpha}$. Making a wrong evaluation decreases the expert's reputation from α to the lower level $\underline{\alpha}$. In the rest of the paper we denote $(\overline{\alpha} - \underline{\alpha})$ as the reputational reward of being recognized as a good expert. This allows us to write conditions (4) and (5) in the following way:

$$k\left(\widehat{\theta}_{\alpha,m_h} - \widehat{\theta}_{\alpha,m_l}\right) \leq (1-k)(\overline{\alpha} - \underline{\alpha})\left(1 - 2\Pr(\omega = h|s_l)\right), \tag{6}$$

$$k\left(\widehat{\theta}_{\alpha,m_h} - \widehat{\theta}_{\alpha,m_l}\right) \geq (1-k)(\overline{\alpha} - \underline{\alpha})\left(1 - 2\Pr(\omega = h|s_h)\right). \tag{7}$$

[15] We show this result in the Appendix.

For each of the above conditions, we refer to the left hand side as the benefit of providing a high message, and to the right hand side as the expected reputational gain of sending a low message. Notice that the right hand side of (6) represents the expected reputational gain of truthtelling when receiving a low signal, while the right hand side of (7) represents the expected reputational gain of misreporting when receiving a high signal.

It is straightforward to show that when reputation does not play any role (i.e., when $k = 1$), condition (6) is never satisfied and a truthtelling equilibrium never exists. We now establish that whenever the expert is concerned about reputation some information can be transmitted.

Lemma 1 *The most informative equilibrium is such that for any* $k \in [0, 1)$*, there always exists a non-empty interval* $[\underline{\theta}, \overline{\theta}]$ *with* $0 < \underline{\theta} < \overline{\theta} < 1$*, such that for* $\theta \in [\underline{\theta}, \overline{\theta}]$ *the equilibrium is separating (i.e. fully revealing), and for* $\theta \notin [\underline{\theta}, \overline{\theta}]$ *the equilibrium is pooling (i.e. uninformative)*

(Proof: see Appendix)

When θ is relatively extreme, the expert believes that any contrarian signal he receives is likely to be incorrect. Being afraid that ex-post incorrect messages may negatively affect his reputation, the expert disregards his private information and reports the signal that is more likely to be correct ex-post. This is the conservative behavior highlighted by Ottaviani and Sorensen (2001, 2006) for the case in which the expert does not have any partisan bias and is *solely* concerned about his reputation.

There is a simple reason why this behavior of the expert's persists in our context with conflicts of interest. When θ is very low (high), the decision maker expects the state to be l (h) regardless of the message sent by the expert. As a result, the net gain from inflating the beliefs of the the decision maker by sending m_h instead of a m_l is very small (i.e., the LHSs of conditions (6) and (7) are close to zero) and the choice of the expert is mainly driven by reputational concerns.

The previous finding highlights that when public opinion is polarized, conflicts of interest are innocuous. No matter how strong is the bias of an expert (i.e. how large is k), his incentives to misreport are small since his ability to sway the decision maker's beliefs in a desired direction is limited. In these cases, inefficiencies in information revelation may arise only because of reputational concerns.

As we show in sections 4 and 5, the presence of conflicts of interest is not innocuous with respect to other relevant dimensions of the problem that we are analyzing.

4 Information Quality and its Effect on Efficiency

In this section, we examine how a variation in the expected quality of the expert's information (i.e., $\alpha p + (1-\alpha)z$) affects informational efficiency. We first study how, *ceteris paribus*, a change in the prior level of reputation α affects the size of the truthtelling region, which we measure with the difference $\overline{\theta} - \underline{\theta}$. We then perform the same exercise but focusing on changes in the difference between the precision of the signals, $p - z$. With a slight abuse of terminology, we refer to any increase (decrease) in $\overline{\theta} - \underline{\theta}$ as to an increase (decrease) in informational efficiency.

Significantly different results arise in the presence of conflicts of interest as opposed to the case in which conflicts are absent. It is useful to consider the case of no conflicts of interest ($k = 0$) as a benchmark case.

Remark 1 *In the absence of conflicts of interest, the truthtelling region* $\left[\underline{\theta}, \overline{\theta}\right]$ *is symmetrically centered around* $\theta = \frac{1}{2}$. *Ceteris paribus, the truthtelling region expands monotonically as either* α, p *or* z *increase.*

(Proof: see Appendix)

Thus, in the absence of conflicts of interest an increase in initial reputation (i.e., an increase in α) or in the precision of the signals (i.e. an increase in either p or z) has a positive effect on informational efficiency. In what follows we study what happens in the presence of conflicts of interest ($k \in (0,1)$).

4.1 Variations in Prior Reputation (α)

In order to understand how variations in α affect informational efficiency when conflicts of interest exist, it is useful to consider the truthtelling conditions (6) and (7) and analyze how α affects the two components $\widehat{\theta}_{\alpha,m_h} - \widehat{\theta}_{\alpha,m_l}$ and $\overline{\alpha} - \underline{\alpha}$.

Remark 2 *(i) The net benefit of sending a high report, $\widehat{\theta}_{\alpha,m_h} - \widehat{\theta}_{\alpha,m_l}$, is increasing in the level of prior reputation α; (ii) The net reputational reward of being recognized as a good expert, $\overline{\alpha} - \underline{\alpha}$, is strictly concave in α, and progressively shrinks to zero as α approaches either zero or one.*

(Proofs: see Appendix)

The intuition behind part (i) of remark 2 is that an expert with a higher level of reputation is expected to have more accurate signals. As a consequence, in any informative equilibrium, the messages of this expert will have a greater impact on the decision maker's beliefs about the state. As for part (ii) of remark 2, the strict concavity of the reputational reward is explained by the fact that the scope for reputation acquisition is greater, the higher the uncertainty about the expert's ability and becomes negligible when the uncertainty about ability is very low.[16] The interaction between these two effects determines the way in which a variation in α affects informational efficiency. As α increases *above* a certain threshold, the net reputational reward of being recognized as a good expert (i.e. $\overline{\alpha}-\underline{\alpha}$) starts to shrink, while the net benefit of sending a high message (i.e. $\widehat{\theta}_{\alpha,m_h}-\widehat{\theta}_{\alpha,m_l}$) keeps growing larger. Eventually, the incentives of the expert to sway the beliefs of decision makers in favor of the high state prevail, thereby reducing informational efficiency. This effect intensifies as α approaches one,

[16] This is consistent with Holmstrom (1999), in which the reputational incentives for an agent to provide effort are positively related to the level of uncertainty on the agent's ability.

in which case the truthtelling region becomes an empty set. A similar reasoning suggests that an increase in α has a positive effect on informational efficiency whenever the initial level of reputation is *below* a certain threshold. The following proposition formalizes these findings.

Proposition 1 *There always exist* $0 < \alpha^* < \alpha^{**} < 1$ *such that an improvement in initial reputation* α *increases informational efficiency if* $\alpha \in (0, \alpha^*)$*, and decreases informational efficiency if* $\alpha \in (\alpha^{**}, 1)$*.*

(Proof: see Appendix)

Proposition 1 suggests that in the presence of conflicts of interest a variation in prior reputation has a non-monotonic effect on informational efficiency. As Figure 1 (left panel) shows, a further increase in prior reputation above a certain threshold makes the truthtelling region shrink. Furthermore, there is little to no scope for information revelation at extreme values of reputation despite the fact that both types of experts have informative signals. These results are in sharp contrast with the case of no conflicts of interest (Figure 1, right panel), where informational efficiency increases monotonically in the level of prior reputation. In the absence of conflicts of interest, if an expert with a certain level of reputation truthfully reveals his information, then any expert with a higher level of reputation will also report truthfully. When conflicts exist, this is not necessarily true. Ceteris paribus, the distorsionary incentives provided by reputation may induce a more reputable expert to lie when a less reputable expert would not. This suggests that in the presence of conflicts of interest, decision makers are not necessarily better off when they consult more reputable experts. In fact, numerical analysis reveals that as the bias becomes more severe (i.e., as k increases), the value of reputation that maximizes informational efficiency diminishes (Fig. 2). It may therefore be optimal for decision makers to consult less reputable experts precisely when incentives are less aligned.

4.2 Variations in Signals' Informativeness

In this section, we examine how variations in the difference between the precision of the signals (i.e., $p - z$) affect informational efficiency. The following proposition summarizes the main finding.

Proposition 2 *Holding p fixed, there always exists a level of z above which an increase in z reduces informational efficiency.*

(*Proofs: see Appendix)*

Notice that, holding p fixed, an increment in z always leads to an increase in the average informativeness of the expert's signals (i.e., $\alpha p + (1-\alpha)z$). Thus, proposition 2 highlights a result whereby informational efficiency may suffer from an improvement in the accuracy of the expert's information. Indeed, as the signal precision of the worst type improves, the difference $\widehat{\theta}_{\alpha,m_h} - \widehat{\theta}_{\alpha,m_l}$ increases. This occurs because the decision maker expects the report of the expert to be more informative on average. At the same time, as z approaches p, the reputational reward of being recognized as a good expert decreases since the difference between good and bad types shrinks. Thus, as the abilities of the two types converge, the amount of information revealed tends to zero.[17] An implication of this result is that the coexistence of experts with different abilities is key to having higher levels of informational efficiency. Again, this is in contrast with the case of no conflicts of interest where an increase in z has an unambiguously positive effect on informational efficiency. In this case, informational efficiency is indeed maximized exactly when $z \to p$.

[17] We also perform the exercise of keeping z fixed and letting p vary. As expected, informational efficiency is increasing in p. Overall, these results are consistent with the idea that the gap in the abilities of the experts plays a key role along with the accuracy of the experts' information.

5 Conflicts of interest and Public Pessimism

We conclude our analysys by examining how variations in k affect the truthtelling thresholds $\underline{\theta}$ and $\overline{\theta}$. The following proposition highlights that when the prior on the state of the world is rather pessimistic, experts with a greater conflict of interest may be more likely to tell the truth than experts with a lower bias.

Proposition 3 *As conflicts of interest increase, the truthtelling region shifts downwards (i.e. both $\underline{\theta}$ and $\overline{\theta}$ are decreasing in k)*

(*Proof: see Appendix)*

We know from remark 1 that in the absence of a bias (i.e., $k = 0$), the truthtelling region is centered around $\theta = \frac{1}{2}$. Proposition 3 suggests that as conflicts of interest become more severe, the truthtelling region progressively shifts towards values of the prior on the state of the world that are closer to zero. The reason for this result is straightforward. Suppose that $k = k'$ and that the expert has received S_l. We know that in this case the expert reports the low message for all values of θ in the region $\left(0, \overline{\theta}_{k'}\right]$, with $\overline{\theta}_{k'}$ denoting the threshold value of θ at which the expert is indifferent between the high and the low message. If k increases above k', the optimistic bias of the expert increases too, and the expert's indifference at $\overline{\theta}_{k'}$ is broken in favour of the high message. Similarly, if the expert has received S_h, he reports the high message for values of θ in the region $[\underline{\theta}_{k'}, 1)$, where $\underline{\theta}_{k'}$ denotes the threshold value of θ at which the expert is indifferent between the high and the low message. Again, as k increases and the expert's optimistic bias becomes stronger, the indifference at $\underline{\theta}_{k'}$ will be broken in favor of the high message.

In general, the previous analysis suggests that when an expert is biased, truthful revelation occurs when public information is rather contrary to the state towards which the expert wishes to sway public opinion. Therefore, in situations in which public information is polarized around a certain belief, consulting an expert with a contrary bias may be more valuable than consulting an unbiased expert.

6 Conclusion

Conflicts of interest are relevant in many economic settings in which an expert with privileged information is called upon to provide information to an uninformed decision maker. In this paper, we focused on the trade-off that biased experts typically face between the short-term benefits of providing a biased report and the long-term rewards of establishing a reputation for having accurate information.

We find that the interaction between reputation and conflicts of interest plays an important role in shaping the incentives of experts. Reputation for ability allows for some information transmission even when decision makers know that an expert is biased. However, reputation has a non-monotonic effect on information transmission. In particular, when the expert's reputation rises above a critical threshold, the expert is more likely to misreport. We found that this critical threshold varies with the degree of conflicts of interest. Specifically, highly reputable experts are more likely to misreport exactly when conflicts of interest are more severe. This occurs because an expert with a higher reputation is more credible. Accordingly, his report has a greater impact on the beliefs of the decision maker. Eventually, the greater is the bias of an expert, the greater his incentives to cash in on reputation by swaying the beliefs of decision makers in the desired direction.

A suggested direction for future research is to gather further insight on the role of both ability and preferences in jointly determining the reputation of experts. In particular, a relevant question involves understanding in which circumstances decision makers may be better off consulting experts that have similar preferences, rather than those that have more accurate information. Capturing how these elements may affect informational efficiency through the reputational channel represents an open issue.

Appendix

Expert's Posterior Beliefs.

$$\Pr(\omega = h|s_h) = \frac{\theta(\alpha p + (1-\alpha)z)}{\theta(\alpha p + (1-\alpha)z) + (1-\theta)(\alpha(1-p) + (1-\alpha)(1-z))};$$
$$\Pr(\omega = l|s_h) = 1 - \Pr(\omega = h|s_h);$$
$$\Pr(\omega = h|s_l) = \frac{\theta(\alpha(1-p) + (1-\alpha)(1-z))}{\theta(\alpha(1-p) + (1-\alpha)(1-z)) + (1-\theta)(\alpha p + (1-\alpha)z)};$$
$$\Pr(\omega = l|s_l) = 1 - \Pr(\omega = h|s_l).$$

■

Posterior Reputations under Truthtelling. In a truthtelling equilibrium the expert reports the signal he has observed. Therefore:

$$\widehat{\alpha}_{\omega,m_j} \equiv \Pr(t = g|\omega, m_j) = \begin{cases} \frac{\alpha p}{\alpha p+(1-\alpha)z} & \text{for } (\omega = h,\ j = h), (\omega = l,\ j = l). \\ \frac{\alpha(1-p)}{\alpha(1-p)+(1-\alpha)(1-z)} & \text{for } (\omega = h,\ j = l), (\omega = l,\ j = h). \end{cases}$$

Let $\overline{\alpha} \equiv \frac{\alpha p}{\alpha p+(1-\alpha)z}$ and $\underline{\alpha} \equiv \frac{\alpha(1-p)}{\alpha(1-p)+(1-\alpha)(1-z)}$. Then for $\alpha \in (0,1)$, $p \in \left(\frac{1}{2}, 1\right)$ and $z \in \left[\frac{1}{2}, p\right)$:

$$\overline{\alpha} - \underline{\alpha} = \frac{\alpha p}{\alpha p + (1-\alpha)z} - \frac{\alpha(1-p)}{\alpha(1-p) + (1-\alpha)(1-z)} = \frac{\alpha(1-\alpha)(p-z)}{(1-\alpha(p-z)-z)(\alpha(p-z)+z)} > 0.$$

■

Proof of Lemma 1. Consider the two conditions for truthtelling:

$$k[\widehat{\theta}_{\alpha,m_h} - \widehat{\theta}_{\alpha,m_l}] \leq (1-k)(\overline{\alpha} - \underline{\alpha})\left[1 - 2\Pr(\omega = h|s_l)\right], \tag{A1}$$
$$k[\widehat{\theta}_{\alpha,m_h} - \widehat{\theta}_{\alpha,m_l}] \geq (1-k)(\overline{\alpha} - \underline{\alpha})\left[1 - 2\Pr(\omega = h|s_h)\right]. \tag{A2}$$

To prove Lemma 1, we first establish and prove the following two results.

Result 1: *In a truthtelling equilibrium, the benefit of sending a high message,* $k\left(\widehat{\theta}_{\alpha,m_h}-\widehat{\theta}_{\alpha,m_l}\right)$ *satisfies the following properties: a) it is strictly positive for* $\theta\in(0,1)$ *and equal to zero for* $\theta=0,1$*; b) it is strictly concave in* θ *with a maximum at* $\theta=\frac{1}{2}$.

Since $k\in[0,1]$, we can analyze $f(\theta)\equiv\widehat{\theta}_{\alpha,m_h}-\widehat{\theta}_{\alpha,m_l}$. In a truthtelling equilibrium the expert reports the signal he has observed. Therefore:

$$\widehat{\theta}_{\alpha,m_j}\equiv\Pr(\omega=h|m_j)=\Pr(\omega=h\mid s_j)=\begin{cases}\frac{\theta(\alpha p+(1-\alpha)z)}{\theta(\alpha p+(1-\alpha)z)+(1-\theta)(\alpha(1-p)+(1-\alpha)(1-z))} & \text{for } j=h\\ \frac{\theta(\alpha(1-p)+(1-\alpha)(1-z))}{\theta(\alpha(1-p)+(1-\alpha)(1-z))+(1-\theta)(\alpha p+(1-\alpha)z)} & \text{for } j=l\end{cases}.$$

With a bit of algebra we obtain:

$$f(\theta)\equiv\widehat{\theta}_{\alpha,m_h}-\widehat{\theta}_{\alpha,m_l}=$$
$$=\frac{\theta(-1+\theta)\left(-1+2\left(\alpha(p-z)+z\right)\right)}{\left(\theta\left(2\left(\alpha(p-z)+z\right)-1\right)-\left(\alpha(p-z)+z\right)\right)\left(1+\theta\left(2\left(\alpha(p-z)+z\right)-1\right)\right)-\left(\alpha(p-z)+z\right)}.$$

Let $q\equiv\alpha(p-z)+z$. Then, $f(\theta)=-\frac{\theta(1-\theta)(2q-1)}{(2q\theta-\theta-q)(1+2q\theta-\theta-q)}$. Notice that for $\alpha\in(0,1)$, $p\in\left(\frac{1}{2},1\right)$ and $z\in\left[\frac{1}{2},p\right)$, we have that $\frac{1}{2}<q<1$. Then:

$f(\theta)>0$ for $0<\theta<1$;

$f(\theta)=0$ for $\theta=0,1$;

$$\frac{\partial f(\theta)}{\partial\theta}=-\frac{q(1-q)(2q-1)(2\theta-1)}{\left(2q\theta-\theta-q\right)^2\left(1+2q\theta-\theta-q\right)^2}\begin{cases}>0 \textit{ for } 0<\theta<\frac{1}{2}\\ =0 \textit{ for } \theta=\frac{1}{2}\\ <0 \textit{ for } \frac{1}{2}<\theta<1\end{cases};$$

$$\frac{\partial^2 f(\theta)}{\partial\theta^2}=2q(1-q)(2q-1)\left(\frac{1}{(2q\theta-\theta-q)^3}-\frac{1}{(1+2q\theta-\theta-q)^3}\right)<0,\ \text{for } 0<\theta<1.$$

Result 2: *In a truthtelling equilibrium, the expected reputational gain of sending the low message,* $(1-k)(\overline{\alpha}-\underline{\alpha})\left(1-2\Pr\left(\omega=h|s_i\right)\right)$ *satisfies the following properties: a) it is positive at* $\theta=0$ *and negative at* $\theta=1$ *for* $i=h,l$*; b) it is strictly decreasing in* θ *for*

$i = h, l$; it is strictly concave in θ for $i = l$ and strictly convex in θ for $i = h$.

Let $g(\theta) \equiv (1-k)(\overline{\alpha}-\underline{\alpha})(1-2\Pr(\omega = h|s_l))$ and $v(\theta) \equiv (1-k)(\overline{\alpha}-\underline{\alpha})(1-2\Pr(\omega = h|s_h))$. Using the values of $\overline{\alpha}$, $\underline{\alpha}$, $\Pr(\omega = h|s_l)$ and $\Pr(\omega = h|s_h)$ we obtain:

$$g(\theta) = \frac{(1-k)(1-\alpha)\alpha(p-z)(-\theta+\alpha(p-z)+z)}{(-1+\alpha(p-z)+z)(\alpha(p-z)+z)(\alpha(-1+2\theta)(p-z)-z+\theta(-1+2z))} \text{ (RHS of (6))},$$
$$v(\theta) = \frac{(1-k)\alpha(1-\alpha)(p-z)(-1+\theta+\alpha(p-z)+z)}{(-1+\alpha(p-z)+z)(\alpha(p-z)+z)(1+\alpha(-1+2\theta)(p-z)-z+\theta(-1+2z))} \text{ (RHS of (7))}.$$

Let $q \equiv \alpha(p-z)+z$. Then, $g(\theta) = \frac{\alpha(p-q)(\theta-q)}{q(1-q)(2q\theta-\theta-q)}$ and $v(\theta) = \frac{\alpha(p-q)(1-\theta-q)}{q(1-q)(2\theta q-\theta-q+1))}$. Notice that for $\alpha \in (0,1)$, $p \in \left(\frac{1}{2}, 1\right)$ and $z \in \left[\frac{1}{2}, p\right)$, we have that $\frac{1}{2} < z < q < p < 1$. Then:

$$g(\theta) \begin{cases} > 0 \textit{ for } 0 < \theta < q \\ = 0 \textit{ for } \theta = q \\ < 0 \textit{ for } q < \theta < 1 \end{cases};$$
$$g(0) = \frac{\alpha(p-q)}{q(1-q)} > 0,\ g(1) = -\frac{\alpha(p-q)}{q(1-q)} < 0;$$
$$\frac{\partial g(\theta)}{\partial\theta} = -\frac{2\alpha(p-q)}{(q+\theta-2q\theta)^2} < 0 \textit{ for } 0 < \theta < 1;$$
$$\frac{\partial g^2(\theta)}{\partial\theta^2} = -\frac{4\alpha(p-q)(2q-1)}{(q+\theta-2q\theta)} < 0 \textit{ for } 0 < \theta < 1;$$

$$v(\theta) \begin{cases} > 0 \textit{ for } 0 < \theta < 1-q \\ = 0 \textit{ for } \theta = 1-q \\ < 0 \textit{ for } 1-q < \theta < 1 \end{cases};$$
$$v(0) = \frac{\alpha(p-q)}{q(1-q)} > 0,\ v(1) = -\frac{\alpha(p-q)}{q(1-q)} < 0;$$
$$\frac{\partial v(\theta)}{\partial\theta} = -\frac{2\alpha(p-q)}{(-1+q+\theta-2q\theta)^2} < 0 \textit{ for } 0 < \theta < 1;$$
$$\frac{\partial v^2(\theta)}{\partial\theta^2} = \frac{4\alpha(p-q)(2q-1)}{(1-q-\theta+2q\theta)^3} > 0 \textit{ for } 0 < \theta < 1;$$

$$g(\theta) - v(\theta) = \frac{2\alpha(p-q)(2q-1)(1-\theta)\theta}{q(1-q)(1-q-\theta+2q\theta)(q+\theta-2q\theta))} > 0, \text{ for } 0 < \theta < 1.$$

We now prove Lemma 1. We first prove that for every value of $\alpha \in (0,1)$, $k \in [0,1)$, $p \in \left(\frac{1}{2}, 1\right)$ and $z \in \left[\frac{1}{2}, p\right)$, there exist $\underline{\theta} \in [0,1]$ and $\overline{\theta} \in [0,1]$ such that for $\theta \in \left[\underline{\theta}, \overline{\theta}\right]$ conditions (A1) and (A2) are satisfied simultaneously. Consider condition (A1) first. Using Results 1 and 2, we can write (A1) as follows:

$$-\frac{k\theta(1-\theta)(2q-1)}{(2q\theta-\theta-q)(1+2q\theta-\theta-q)} \leq \frac{(1-k)\alpha(p-q)(\theta-q)}{(1-q)q(2q\theta-\theta-q)}.$$

Notice that $\frac{1}{2} \leq z < q < p < 1$. Thus, for $\theta \in (0,1)$, $2q\theta - \theta - q < 0$ and (A1) is equivalent to:

$$\frac{k\theta(1-\theta)(2q-1)}{1+2q\theta-\theta-q} \leq -\frac{(1-k)\alpha(p-q)(\theta-q)}{(1-q)q}. \tag{A3}$$

Finally, let $h(\theta) = -\frac{k\theta(1-\theta)(2q-1)}{2q\theta-\theta-q}$ and $r(\theta) = \frac{(1-k)\alpha(p-q)(\theta-q)}{(1-q)q}$, and notice that:

a) $r(0) > h(0) = 0$, $r(1) < h(1) = 0$

b) $r(\theta)$ is a negatively sloped straight line.

c) $h(\theta)$ is non-negative, continuous, and strictly concave for $\theta \in (0,1)$.

Properties a), b) and c) imply that there exists a unique $\overline{\theta} \in (0,1)$ such that for any $\theta < \overline{\theta}$ (A3) (and therefore (A1)) are satisfied.

Focusing on condition (A2) and following the same steps above, we can prove the existence and uniqueness of a $\underline{\theta} \in (0,1)$ such that, for any $\theta > \underline{\theta}$, (A2) is satisfied. From Result 2 we know that for $\theta \in (0,1)$ the RHS of condition (A1) is strictly greater than the RHS of condition (A2). This result, together with the uniqueness of $\underline{\theta}$ and $\overline{\theta}$ implies that $\overline{\theta} > \underline{\theta}$. Therefore, (A1) and (A2) are simultaneously satisfied for $\theta \in \left[\underline{\theta}, \overline{\theta}\right]$.

Finally, notice that a babbling equilibrium where the expert sends m_h with probability π and m_l with probability $1-\pi$ irrespectively of the signal observed always exists. In this case all messages are taken to be meaningless and ignored: $\widehat{\theta}_{\alpha,m_j} = \theta$ for any $i = h, l$, and $\widehat{\alpha}_{\omega,m_j} = \alpha$ for any $\omega = h, l$ and $j = h, l$, making the expert indifferent between the two messages. ■

Proof of Remark 1. When $k = 0$, condition (A3) boils down to $0 \leq \alpha(p-q)(\theta-q)$. Thus, when $k = 0$, the upper bound of the truthtelling region is $\overline{\theta} = q \equiv \alpha p + (1-\alpha)z$. Similarly, from condition (A2), one can show that when $k = 0$ the lower bound of the truthtelling region is $\underline{\theta} = 1 - (\alpha p + (1-\alpha)z)$.

Proof of Remark 2. Let $q = \alpha(p-z) + z$, where $z < q < p$. Notice that: ■

(i) $\frac{\partial\left(\widehat{\theta}_{\alpha,m_h} - \widehat{\theta}_{\alpha,m_l}\right)}{\partial\alpha} = \theta(1-\theta)(p-z)\left(\frac{1}{(q+\theta(1-2q))^2} + \frac{1}{(1-q-\theta(1-2q))^2}\right) > 0$ for any $\alpha \in (0,1)$.

(ii) $\frac{\partial(\overline{\alpha}-\underline{\alpha})}{\partial\alpha} = \frac{(p-z)\left(\alpha^2(p-1)p+(\alpha-1)^2z-(\alpha-1)^2z^2\right)}{(q-1)^2q^2}$; Notice that: $\frac{\partial(\overline{\alpha}-\underline{\alpha})}{\partial\alpha} = 0 \Leftrightarrow \alpha_0 = \frac{z-z^2-\sqrt{pz-p^2z-pz^2+p^2z^2}}{p^2-p+z-z^2}$, $\alpha_1 = \frac{z-z^2+\sqrt{pz-p^2z-pz^2+p^2z^2}}{p^2-p+z-z^2}$, where $\alpha_1 < 0 < \alpha_0 < 1$.

(iii) $\frac{\partial^2(\overline{\alpha}-\underline{\alpha})}{\partial\alpha^2} = 2(p-z)\left(-\frac{(1-p)(1-z)}{(1-q)^3} - \frac{pz}{q^3}\right) < 0$ for $\alpha \in (0,1)$.

Therefore, for $\alpha \in (0,1)$, $\overline{\alpha} - \underline{\alpha}$ is strictly concave with a maximum at $\alpha = \alpha_0$. Also, since $\overline{\alpha} \equiv \frac{\alpha p}{\alpha p+(1-\alpha)z}$ and $\underline{\alpha} \equiv \frac{\alpha(1-p)}{\alpha(1-p)+(1-\alpha)(1-z)}$, we have that $\overline{\alpha} - \underline{\alpha} \to 0$ as $\alpha \to 0$. ■

Proof of Proposition 1. Consider condition (A1) and notice that: (i) For $\alpha \to 0$, $LHS_1 \to \frac{k\theta(2z-1)(1-\theta)}{(2z\theta-\theta-z)(2z\theta-\theta-z+1)}$ and $RHS_1 \to 0$; thus, for $\alpha \to 0$, $\overline{\theta} \to 0$; (ii) For $\alpha \to 1$, $LHS_1 \to \frac{k\theta(2p-1)(1-\theta)}{(2p\theta-\theta-p)(2p\theta-\theta-p+1)}$ and $RHS_1 \to 0$; thus, for $\alpha \to 1 : \overline{\theta} \to 0$. A similar argument applies to condition (A2) to show that: (iii) For $\alpha \to 0$, $\underline{\theta} \to 0$; (iv) For $\alpha \to 1$, $\underline{\theta} \to 0$. Now, (i) and (ii) imply that for $\alpha \to 0$, $\overline{\theta} - \underline{\theta} \to 0$; (iii) and (iv) imply that for $\alpha \to 1$, $\overline{\theta} - \underline{\theta} \to 0$. Since $\overline{\theta} - \underline{\theta}$ is positive for any value of $\alpha \in (0,1)$, by continuity there exist a value of $\alpha \in (0,1)$ below which $\overline{\theta} - \underline{\theta}$ is increasing in α, and a value of $\alpha \in (0,1)$ above which $\overline{\theta} - \underline{\theta}$ is decreasing in α. ■

Proof of Proposition 2. Consider conditions (A1) and (A2).Notice that for $z \to p$: (i) $LHS_1 \to \frac{k\theta(1-\theta)(2p-1)}{(2p\theta-\theta-p)(2p\theta-\theta-p+1)}$ and $RHS_1 \to 0$, which implies that $\overline{\theta} \to 0$; (ii) $LHS_2 \to \frac{k\theta(1-\theta)(2p-1)}{(2p\theta-\theta-p)(2p\theta-\theta-p+1)}$ and $RHS_2 \to 0$, which implies that $\underline{\theta} \to 0$. From (i) and (ii) it follows that for $z \to p$, $\overline{\theta} - \underline{\theta} \to 0$. Since $\overline{\theta} - \underline{\theta}$ is positive for any value of $z \in (0,p)$, by continuity there exist a value of $z \in (0,1)$ above which $\overline{\theta} - \underline{\theta}$ is decreasing in z. ■

Proof of Proposition 3. To proof proposition 3, first note that:

$$\text{In condition (A1), } \left.\frac{\partial RHS}{\partial \theta}\right|_{\theta=\overline{\theta}} > \left.\frac{\partial LHS}{\partial \theta}\right|_{\theta=\overline{\theta}}; \tag{8}$$

$$\text{In condition (A2), } \left.\frac{\partial RHS}{\partial \theta}\right|_{\theta=\underline{\theta}} > \left.\frac{\partial LHS}{\partial \theta}\right|_{\theta=\underline{\theta}}; \tag{9}$$

This result is a straightforward consequence of the uniqueness of $\overline{\theta}$ and $\underline{\theta}$, together with the properties highlighted in the Result 1 and Result 2 in the proof of Lemma 1. In words, (8) says that the RHS of (A1) always intersects the LHS from above. (9) says that the same is true for condition (A2).

Now consider condition (A1). We know from Result 1 in the proof of Lemma 1 that the LHS of (A1) is strictly positive for any $\theta \in (0,1)$. This implies that at $\theta = \overline{\theta}$, $RHS = LHS > 0$. Further, we know from Result 2 in the proof of Lemma 1 that the RHS of (A1) is strictly decreasing in θ for any $\theta \in (0,1)$, and equal to zero at $\theta = q$. Therefore, it must be that $\overline{\theta} < q$. Having established this result, it is straightforward to show that for any $\theta \in (0,q)$, $\frac{\partial LHS}{\partial k} = -\frac{\theta(1-\theta)(2q-1)}{(2q\theta-\theta-q)(1+2q\theta-\theta-q)} > 0$ and $\frac{\partial RHS}{\partial k} = -\frac{\alpha(p-q)(\theta-q)}{(1-q)q(2q\theta-\theta-q)} < 0$. This result, together with (8) and (9) implies that $\overline{\theta}$ is decreasing in k. The same reasoning applies to condition (A2) to show that $\underline{\theta}$ is decreasing in k. ■

References

[1] Alesina, A., Roubini, N., and Cohen, G. D., 1997. Political Cycles and the Macroeconomy. Cambridge, MA: MIT Press.

[2] Banks, J. S., Weingast, B. R., 1989. The Political Control of Bureaucracies under Asymmetric Information. *American Journal of Political Science*, 36, 509–24.

[3] Barber, B., Lehavy, R., McNichols, M., and Trueman, B., 2006. Buys, holds, and sells: The distribution of investment banks' stock ratings and the implications for the profitability of analysts' recommendations. *Journal of Accounting and Economics*, 41, 87–117.

[4] Barber, B., Lehavy, R., and Trueman, B., 2007. Comparing the stock recommendation performance of investment banks and independent research firms. *Journal of Financial Economics*, 85, 490-517.

[5] Battigalli, P., Dufwenberg, M., 2009. Dynamic Psychological Games. *Journal of Economic Theory*, 144, 1-35.

[6] Benabou, R., Laroque, G., 1992. Using Privileged Information to Manipulate Markets: Insiders, Gurus and Credibility. *Quarterly Journal of Economics*, 107, 921-958.

[7] Bendor, J., Taylor, S., and Van Gaalen, R., 1985. Bureaucratic Expertise versus Legislative Authority: A Model of Deception and Monitoring in Budgeting. *American Political Science Review*, 70, 1041–1060.

[8] Bourjade, S., Jullien, B., 2011. The roles of reputation and transparency on the behaviour of biased experts. *Rand Journal of Economics*, 42, 575-594.

[9] Carlsen, F., 1999. Inflation and Elections: Theory and Evidence for Six OECD Economies. *Economic Inquiry*, 37, 120–135.

[10] Carpenter, D. P., 2001. The Forging of Bureaucratic Autonomy: Reputations, Networks, and Policy Innovation in Executive Agencies, 1862–1928, Princeton, NJ: Princeton University Press.

[11] Crawford, V. P., Sobel, J., 1982. Strategic Information Transmission. *Econometrica*, 50, 1431-1451.

[12] Ely, J., Valimaki, J., 2003. Bad Reputation, *Quarterly Journal of Economics*, 118, 785-814.

[13] Fang, L., Yasuda, A., 2009. The Effectiveness of Reputation as a Disciplinary Mechanism in Sell-side Research. *Review of Financial Studies*, 22, 3735-3777.

[14] Gentzkow, M., Shapiro, J. M., 2006. Media Bias and Reputation. *Journal of Political Economy*, 114, 280-316.

[15] Heclo, H., 1975. OMB and the Presidency - The Problem of 'Neutral Competence'. *The Public Interest*, 38, 80–99.

[16] Holmström, B., 1999. Managerial Incentive Problems: A Dynamic Perspective. *Review of Economic Studies*, 66, 169–182.

[17] Hong, H., Kubik, J.D., 2003. Analyzing the Analysts: Career Concerns and Biased Earnings Forecasts. *Journal of Finance*, 58, 313–351.

[18] Michaely, R., Womak, K. L., 1999. Conflict of Interest and the Credibility of Underwriter Analyst Recommendations. *Review of Financial Studies*, 12, 653-686.

[19] Mikhail, M.B., Walther, B.R., and Willis, R.H., 1999. Does forecast accuracy matter to security analysts?. *Accounting Review*, 74, 185–200.

[20] Morgan, J., Stocken, P. C., 2003. An Analysis of Stock Recommendations. *RAND Journal of Economics,* 34, 183-203.

[21] Morris, S., 2001. Political Correctness. *Journal of Political Economy*, 109, 231-265.

[22] Ottaviani, M., Sorensen, P. N., 2001. Information aggregation in debate: who should speak first?. *Journal of Public Economics,* 81, 393-421.

[23] Ottaviani, M., Sorensen, P. N., 2006. Reputational Cheap Talk. *Rand Journal of Economics*, 37, 155-175.

[24] Rourke, F. E., 1992. Responsiveness and Neutral Competence in American Bureaucracy. *Public Administration Review*, 52, 539–546.

[25] Scharfstein, D. S., Stein, J. C., 1990. Herd Behavior and Investment. *American Economic Review*, 80, 465-479.

[26] Sobel, J., 1985. A Theory of Credibility. *Review of Economic Studies*, 52, 557–573.

[27] Trueman, B., 1994. Analyst Forecasts and Herding Behavior. *Review of Financial Studies*, 7, 97-124.

[28] Weatherford, M. S., 1987. How Does Government Performance Influence Political Support?. *Political Behavior*, 9, 5–28.

[29] Wilson, J. Q., 1989. Bureaucracy: What Government Agencies Do and Why They Do It. New York: Basic Books.

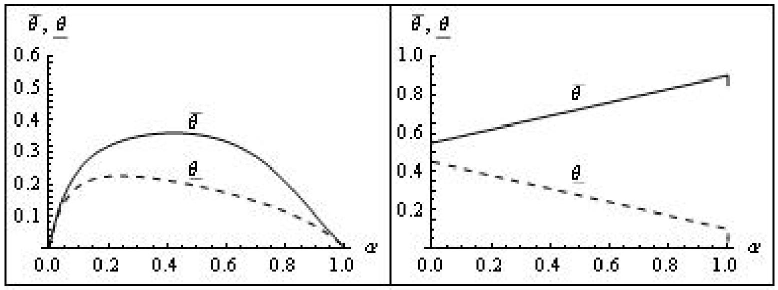

Figure 1: Impact of a variation in α on the size and the position of the truthtelling region in the case of conflicts of interest (left panel) and no conflicts of interest (right panel)

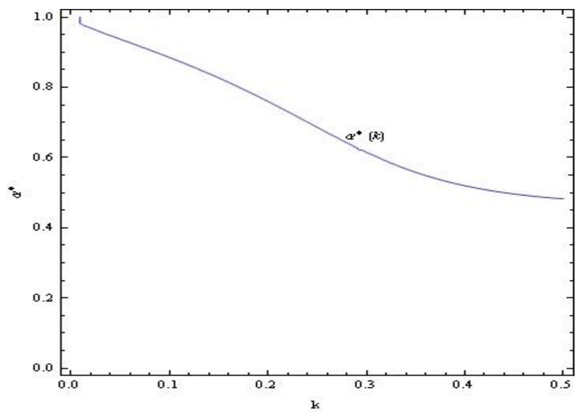

Figure 2: Impact of the bias (k) on the level of reputation that maximizes informational efficiency (α^*).