Good Lies*

Filippo Pavesi[†], Massimo Scotti[‡]

April 2017

Abstract

Decision makers often face uncertainty in relation to the ability and the integrity of their advisors. Whenever this is the case, we show that if an expert is sufficiently concerned about establishing a reputation for being skilled and unbiased, she may truthfully report her private information regarding the decision-relevant state. However, while in a truthtelling equilibrium, the decision maker learns only about the ability of the expert, in an equilibrium with some misreporting the decision maker also learns about the expert's bias. Although truthful behavior allows for more informed current decisions, it may lead to worse sorting outcomes. Therefore, if a decision maker places enough weight on future choices relative to present ones, lying may be welfare improving. Applications of the model include relationships between patients and doctors, managers and consultants, and politicians and policy advisors.

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

JEL Classification: C72, D82, D83

*We would like to thank Emiliano Catonini, Boyan Jovanovic, Navin Kartik, Marco Ottaviani, Andrea Prat, Stephen Morris and Toru Suzuki as well as seminar participants at University of Technology Sydney, University of Verona, University of Wollongong, Bocconi University and conference participants at GAMES 2016 and the Game Theory Society Conference in Stony Brook, NY for helpful comments and suggestions.

†Department of Economics, University of Verona, Via Cantarane 24, 37129 Verona, Italy. Phone: +39 328 3255281; Email: filippo.pavesi@univr.it. Stevens Institute of Technology, School of Business, Castle Point on Hudson, 07030, Hoboken, NJ, USA.

‡Economics Discipline Group, University of Technology, Sydney, PO Box 123 Broadway NSW 2007 Australia. Phone: +61 2 9514 4022; Email: massimo.scotti@uts.edu.au.

1 Introduction

Consider a politician that hires an advisor to make a more informed decision on a specific policy. The politician knows the advisor is a specialist but does not know how well informed the advisor is (*ability*), and whether she is biased in favor of a specific interest group (*integrity*). Suppose the advisor is actually biased and yet provides a recommendation that is genuinely based on her expertise, as an unbiased one would do. As the advisor is a specialist, a truthful recommendation is likely to be correct. This is desirable for the sake of current decisions. However, in deciding whether to consult the specialist again in the future, the politician also requires information about her integrity. Indeed, if the advisor favors a certain industry, she will provide biased recommendations as soon as it is to her advantage. If her behavior is truthful, however, the politician is prevented from learning about the integrity of the advisor because the behavior of a biased expert is indistinguishable from that of an unbiased one (i.e., they both behave truthfully). Now compare the previous situation with a scenario in which, for example, biased advisors tend to ignore what their private information would suggest and prescribe a policy that favors a tax break for a specific industry. In this case, observing such a recommendation may cast doubt on the advisor's integrity. Lying thus reveals evidence about preferences, which remains concealed if behavior is truthful. This knowledge can prove useful in deciding whether to replace the expert or to continue to rely on her services in order receive better advice in the future.

A natural question that arises in this setting is whether facing advisors that sometimes lie in the current period, instead of always honestly reporting their information, may allow politicians to make more informed future decisions. To put it more bluntly: is there scope for good lies? To address this issue, we introduce a model that incorporates the key features of the example described above and identify situations in which some degree of misreporting may be preferable to truthful reporting.

The model we propose is general enough to encompass other settings that involve ongoing relationships between decision makers and experts such as those between patients and doctors, firms and consultants, and investors and financial analysts. In this context, the primary focus of our analysis is on the decision maker's welfare. Specifically, we consider a two-period model of career concerns in which a decision maker chooses a binary action in each period, and her payoff from the action depends on an unknown state of the world. In each period the decision maker can consult an expert that has privileged information about the state, but faces uncertainty about both the ability (i.e., the precision of her information) and the integrity of the advisor (i.e., whether she is biased in favor of a particular course of action). We assume that ability and integrity are independently distributed. The decision maker starts with some prior beliefs about ability and integrity, and updates these beliefs at the end of the first period after observing the recommendation of the expert and the true state of the world. These posterior beliefs are interpreted as the expert's reputations for ability and integrity, and determine how valuable her advice is expected to be in the second period. In particular, these values determine whether the decision maker retains the expert, and if so, how much payment the expert receives for her services in the second period. This, in turn, creates reputational concerns on the part of the expert in the first period.

We show that reputational concerns may induce both biased and unbiased experts to truthfully reveal their information about the state of the world in the current period (*discipline effect*). This is clearly beneficial for the quality of the decision maker's current decisions. The quality of future decisions is instead affected by how much the decision maker learns about the expert's ability and bias (*sorting effect*). In this respect, we note that there is a trade-off between what the decision maker learns about each of these two dimensions. In particular, while truthful reporting allows for sharp learning about the ability of the expert, it nevertheless precludes learning about integrity. Intuitively, this occurs because in a truthtelling equilibrium observing the expert's recommendation is equivalent to observing the expert's information. Hence, the decision maker is in a good position to evaluate the quality of the expert's signal. However, as both biased and unbiased experts behave exactly in the same way (i.e., they both report their information truthfully) and are both as likely to have the same information (i.e., ability and integrity are independent), it is impossible for the decision maker to infer something about the integrity of the expert by simply observing her recommendation. On the contrary, we show that equilibria in which experts only partially reveal their information about the state are such that the reporting strategies of biased and unbiased experts are necessarily different. In these equilibria, while observing a certain recommendation reveals some information about integrity, learning about ability becomes less sharp because the reported recommendation only partially reflects the actual quality of the expert's information.

Our main result shows that equilibria with some misreporting can improve sorting with respect to truthful reporting when they allow for certain specific patterns of learning about ability and integrity. In these cases, decision makers may prefer some misreporting if they are sufficiently concerned about the expected quality of their future decisions, in other words if they have a relative preference for sorting over discipline. We, therefore, prove that although truthtelling equilibria exist, they can be welfare-dominated by equilibria that involve some degree of misreporting. Including two dimensions of reputation thus provides novel results with respect to settings in which there is only one dimension. In these latter cases, either truthtelling equilibria do not exist (as in Morris, 2001 where reputation is only related to preferences and biased advisors prefer actions to be distorted in a particular known direction) or, when they do exist, they always dominate misreporting equilibria (as in Prat, 2005 and Ottaviani and Sorensen, 2006 where reputation is only due to ability).

First, we characterize a class of misreporting equilibria that have the potential to improve sorting with respect to truthtelling. In these equilibria, which we name *Misreporting Biased* (MB), the unbiased expert always truthfully reports her information, the biased expert misreports her information by sometimes recommending the action she favors when her private information would suggest the opposite, and the decision maker retains the expert if and only if her recommendation is ex-post correct. We show that these equilibria improve sorting whenever the prior probability that the expert is well-informed is sufficiently high, and thus there is little scope for learning about ability. This proves that there exist instances in which learning about integrity is relatively more valuable than learning about ability for making future decisions.

Going back to the politician-advisor example, our result suggests that a politician whose current decisions are relatively less important than future ones may prefer a setting in which biased advisors tend to provide advice guided by their conflicts of interest. This will lead the politician to make more mistakes in the present but will allow her to better discriminate between biased and unbiased experts. This is so because equilibrium behavior implies that an advisor that suggests a policy that fails to deliver the desired results will be replaced, and advisors that provide biased suggestions end up making mistakes more often. Whenever the skills of the advisor are less of an issue, as in the case in which the politician can hire a new expert by picking from a pool of experienced policy analysts, a setting with some lying improves sorting with respect to a setting with truthful behavior.

We then characterize the class of equilibria in which the unbiased expert misreports and analyze whether these equilibria have the potential to improve welfare with respect to truthtelling. Like MB equilibria, this class also displays the feature that the expert's recommendation reveals some information about her integrity, and can be further divided into two subclasses. The first subclass, which we denote *Misreporting Unbiased* (MU), is characterized by the unbiased expert partially revealing her information about the state, the biased expert either truthtelling or partially revealing her information depending on her level of career concerns, and the decision maker retaining the expert if and only if her recommendation is ex-post correct. The second subclass, which we denote *Total Misreporting Unbiased* (TMU), is characterized by the unbiased expert always recommending the action that is least-preferred by the biased expert. In these equilibria, the decision maker adopts a rather conservative strategy: she ignores the ex-post correctness of the recommendations and retains the expert only if the recommended action is the one that is least-preferred by the biased expert.

When we consider MU equilibria in which the unbiased expert misreports and the biased expert truthfully reveals her information, we find that they never improve sorting relative to truthtelling equilibria. This is rather surprising because the reporting strategies of MU equilibria would suggest a pattern of learning about ability and integrity and hence a sorting effect similar to the one we have in MB. In fact, we find that misreporting by the unbiased expert hampers the sorting effect because it diminishes the decision maker's chances of consulting an unbiased expert of high ability in the future. Therefore, even when there is little scope for learning about ability, the sorting effect that comes from the unbiased expert's intention of signaling her integrity is not sufficient to offset the sorting effect associated with truthtelling that derives from greater learning about ability.

The previous result may suggest that when unbiased experts misreport to signal their type, this is never optimal from the decision maker's perspective. However, we show that this is not the case, and that there exists a non-empty set of TMU equilibria that can dominate truthtelling. To illustrate this result, consider a patient-doctor relationship. In this setting, TMU equilibria can be described as follows: a patient consults a doctor with the intention to follow her current advice but to switch to a new doctor in the future if her diagnosis suggests undergoing a specific treatment from which it is well known that a biased physician may directly benefit. In this equilibrium, an unbiased doctor that does not face a conflict of interest will never suggest undergoing treatment even if her diagnosis suggests that this is the best current choice for the patient. On the contrary, a biased doctor will suggest undergoing treatment with a positive probability, because she can profit from carrying out the treatment today even knowing that the patient will not return in the future. This behavior allows the patient to learn something about the physician's integrity. If the patient is more concerned about future consultations, in which the odds of facing serious health issues are higher, this scenario prescribed by TMU equilibria will be preferred to one in which both biased and unbiased doctors provide honest evaluations based on their expertise.

Finally, to provide a more complete picture of our findings, we characterize informative equilibria based on the level of reputational concerns of the expert. We show that truthtelling can be sustained only when the expert's career concerns are sufficiently high. However, we find that TMU equilibria may also exist in this case, thus undermining the potential for truthful reporting to be welfare maximizing. Moreover, when career concerns are mild and truthtelling cannot be supported, there exist misreporting equilibria such as MB, which have the potential to dominate truthtelling in terms of welfare. This suggests that it may not always be optimal for a decision maker to consult experts with high reputational concerns.

Our work builds on the existing literature that studies the effects of reputational concerns within models of expertise. This literature has alternatively focused either on reputation for ability (Scharfstein and Stein, 1990; Trueman, 1994; Holmstrom, 1999; Ottaviani and Sorensen, 2006) or for preferences (Sobel, 1985; Benabou and Laroque,1992; Morris, 2001; Ely and Valimaki, 2003). A contribution of the present paper is to propose a model that incorporates both these sources of reputational concerns.

In particular, Morris (2001) and Ely and Valimaki (2003) highlight how reputational concerns may be self-defeating and therefore useless in aligning incentives. In both papers, reputational concerns lead a good agent to engage in inefficient behavior for signaling purposes. In Morris (2001), when reputational concerns are strong, information revelation completely breaks down and babbling is the only equilibrium. Eli and Valimaki (2003) consider an infinite-horizon principal-agent model, and show that principals anticipate the "bad reputation" effect and hence never hire an agent, thereby leading to the loss of all surplus. Although our focus is different because we concentrate on comparing the welfare properties of different informative equilibria, our model provides some insight on these results. With respect to Morris (2001), we show that as long as there is some uncertainty regarding ability, informative equilibria always exist if experts' reputational concerns are high. This suggests that Morris' result that reputational concerns can be self-defeating when they are too pronounced crucially depends on the existence of a single dimension of uncertainty. Ely and Valimaki (2003) derive their bad reputation result under the assumption that principals are short-run players. In fact, they show that if principals are long-run players, the positive value of reputation can be restored, as principals can internalize the value of learning about the type of the agent. Our model also exploits this learning feature, but in a cheap-talk environment and relying on a finite horizon. In particular, as we consider two dimensions of reputation, in our framework, learning about preferences comes at the cost of learning about ability. We, therefore, focus on comparing which of these two effects dominates in different circumstances.

Our paper is also related to Prat (2005), who studies welfare in a static model of expertise in which the agent bears reputational concerns only for ability, and the principal learns about the ability-type of the agent. We also analyze welfare, but we consider two dimensions of uncertainty and endogenously derive the value of information in a two-period model of reputational cheap talk, in the spirit of Li (2007) and Morris (2001). In particular, while in Prat (2005) the discipline and sorting effects go in the same direction (i.e., equilibria with better discipline also display better sorting), in our setting with two dimensions of reputation, there may be a trade-off between the two.

Another strand of literature that is related to our work is the signaling literature that considers agents that are heterogeneous on two dimensions (Austen-Smith and Fryer, 2005; Esteban and Ray, 2006; Bagwell, 2007; and Frankel and Kartik, 2016). In particular, there is a parallel between our analysis and that of Frankel and Kartik (2016). They show that there is a trade-off between the information that can be revealed on each of two dimensions of uncertainty when only one action is available. In this context, learning about one dimension versus the other depends on the cost of signaling, while in our setting, it depends on the equilibrium communication strategy of the experts. A significant difference with respect to this literature is that we incorporate learning about our two dimensions of heterogeneity (i.e., ability and integrity) in an endogenous expression for the value of information. This allows us to evaluate how learning about each dimension affects the decision maker's welfare.

The remainder of the paper is organized as follows. In Section 2, we introduce the general setup of the model and present a preliminary equilibrium analysis. In Section 3, we introduce the main elements of welfare analysis and illustrate how misreporting equilibria necessarily involve more learning about integrity and less about ability with respect to truthtelling. In Section 4, we characterize the informative equilibria in which the unbiased expert reports truthfully and analyze the welfare properties of these equilibria to illustrate our main results. Section 5 discusses other informative equilibria. In Section 6, we present a complete mapping of all the equilibria providing general welfare results. In Section 7, we discuss the crucial role of reputation for ability, and Section 8 concludes.

2 The Model

There are two periods $t = 1, 2$. In each period, a risk-neutral decision-maker (DM) has to choose an action $a_t \in \{0, 1\}$ and receives a payoff $R_t(a_t, x_t)$ that depends on both a_t and the state of the world $x_t \in \{0, 1\}$ as follows:

$$R_t(a_t, x_t) = \begin{cases} r & \text{if } a_t = 1, x_t = 1 \\ -r & \text{if } a_t = 1, x_t = 0 \\ 0 & \text{if } a_t = 0. \end{cases}$$

where $r > 0$.

We assume that in each period, states $x_t = 0$ and $x_t = 1$ occur with equal probability, and that states x_1 and x_2 are independently distributed.[1] At the moment of choosing a_t, DM does not observe the realization of x_t but can consult an expert who has access to a signal $s_t \in (0, 1)$ that is potentially informative about x_t. The expert observes s_t and then reports a message $m_t \in (0, 1)$ to DM, and is paid a fixed fee w_t for her services.

[1]The assumption of a fair prior is not without a loss of generality. However, the results of the paper hold whenever the prior on the state is not too extreme. A setting with a fair prior represents the situation in which uncertainty about the state is highest, and it is thus more likely that DM seeks the advice of an expert.

The assumption of a fixed fee reflects the contractual incompleteness that is typical of the situations we are modelling, in which both the state of the world and the report of the advisor are observable but not verifiable; thus, contracts cannot be written conditional on reports or on the accuracy of reports.

We can think of DM's decision as the decision to invest ($a_t = 1$) or not invest ($a_t = 0$) in a project or asset whose return is uncertain, and we can think of the expert as a consultant or a financial advisor. However, as we mentioned previously, the model is sufficiently general to represent many situations that involve ongoing relationships between a decision maker and an expert, such as those between patients and doctors, firms and consultants, or politicians and policy advisors. Throughout the paper, we will alternately refer to some of these examples to illustrate our findings.

We assume that there is a finite pool of risk-neutral experts and that DM can consult only one expert per period. Experts differ in their preferences and in their ability. However, DM observes neither the preferences nor the ability of an expert.

Expert's ability. An expert can be either smart (S) or dumb (D). A smart expert receives an informative signal, while a dumb expert receives an uninformative signal as modelled by the following signal technology:

$$\Pr(s_t = x_t \mid x_t, S) = p > \Pr(s_t = x_t \mid x_t, D) = 1/2.$$

As it is customary in models of career concerns, we assume that an expert does not know her own ability.[2] We denote α as the common prior about an expert being smart and $q \equiv \alpha p + (1-\alpha)\frac{1}{2}$ as the ex-ante expected precision of an expert's signal.

Expert's preferences. An expert can be either unbiased (U) or biased (B). While an unbiased expert does not favor any particular action, a biased expert strictly prefers $a_t = 1$. We assume that an expert knows her own preferences and let γ denote the common prior about an expert being unbiased. In the remainder of the paper we will refer to the quality of being unbiased as integrity. We also assume that there is no correlation between ability and integrity, so that unbiased and biased experts have the same chances of being smart.

Payoffs and welfare. We model stage-payoffs as follows. A biased expert gets a stage-

[2]Given our signal structure, the assumption of a fair prior about the state of the world guarantees that an expert does not learn anything about her own ability by observing her own signal.

payoff equal to $w_t + a_t$, where a_t is assumed to be relation-specific. Namely, a biased expert receives a_t in period t if and only if the expert has been hired by DM in period t.[3] An unbiased expert faces no conflict of interest and gets a stage-payoff equal to w_t. Finally, we assume that DM's stage-payoff is equal to $R_t(a_t, x_t)$. This is equivalent to assuming that while an expert seeks to maximize her monetary payoff, DM is only concerned about choosing the best state-contingent action in each period. When we analyze welfare, we thus focus exclusively on the decision maker's utility.

We assume that agents may assign different weights to their stage-payoffs. We let $\delta_E \in (0,1)$ denote the weight that an expert assigns to her future payoff relative to her current payoff. Thus, the total payoff of an unbiased expert and the total payoff of a biased expert respectively read:

$$\begin{aligned} \Pi_U &= (1-\delta_E)w_1 + \delta_E w_2, \\ \Pi_B &= (1-\delta_E)(w_1 + a_1) + \delta_E(w_2 + a_2). \end{aligned}$$

Similarly, we let $\delta_{DM} \in (0,1)$ denote the weight that DM assigns to her future payoff relative to her current payoff. Thus, DM's total payoff reads:

$$\Pi_{DM} = (1-\delta_{DM})R_1(a_1, x_1) + \delta_{DM}R_2(a_2, x_2).$$

Hence, in analyzing welfare, we will focus on the expected value of Π_{DM}.

2.1 Reputations and the Value Function

We model reputations by assuming that at the end of the first period, state x_1 is publicly revealed and that DM uses the realization (m_1, x_1) to update her prior beliefs about the ability and the integrity of the incumbent expert. We respectively denote with $\widehat{\alpha}(m_1, x_1) \equiv \Pr(S \mid m_1, x_1)$ and $\widehat{\gamma}(m_1, x_1) \equiv \Pr(U \mid m_1, x_1)$ DM's posterior beliefs about the ability and the integrity of the incumbent expert. These two beliefs respectively represent the reputations that the incumbent has established at the end of the first period for being smart

[3]This is, for example, the case of a financial analyst who may obtain some side benefits if she persuades an investor to make an investment, or of a doctor that receives a higher compensation if she convinces a patient to undergo surgery.

and for being unbiased. They also summarize what DM has learned about the ability and the integrity of the incumbent after interacting with her. We denote the corresponding update on the expected precision of the incumbent's signal with $\widehat{q}(m_1, x_1) \equiv \widehat{\alpha}(m_1, x_1)p + (1 - \widehat{\alpha}(m_1, x_1))\frac{1}{2}$.

As we will formally see in the next section, both $\widehat{\alpha}(m_1, x_1)$ and $\widehat{\gamma}(m_1, x_1)$ positively affect the value of the incumbent's information in the second period. Intuitively, the smarter and the more unbiased the incumbent is, the more useful her information is. We let $V(m_1, x_1)$ denote the value of the incumbent's information in the second period, and refer to $V(m_1, x_1)$ as the value function. Essentially, $V(m_1, x_1)$ maps the reputation of the incumbent for being unbiased and that for being smart (or equivalently what DM has learned about the incumbent's ability and integrity) into the expected value of the incumbent's information in the second period.

We introduce reputational concerns on the part of experts via two channels. First, we assume that at the beginning of the second period, DM computes $V(m_1, x_1)$ and decides whether to retain the incumbent or replace her with a new expert. In this latter case, the new expert is randomly selected from the original pool of experts. Hence, the value of the information of a new expert is independent from what happened in the first period and depends on the prior beliefs α and γ. We will denote the value function of a new expert with V. As we will see, DM will retain the incumbent whenever $V(m_1, x_1) \geq V$. A second channel of career concerns comes from the fee that is paid to the expert in the second period, w_2. In particular, we assume that w_2 is set equal to the value of the expert's information in the second period. Hence, for the incumbent expert, we have that $w_2 = V(m_1, x_1)$.[4,5] All this implies that the incumbent is concerned about maximizing the value of her reputations $\widehat{\alpha}(m_1, x_1)$ and $\widehat{\gamma}(m_1, x_1)$ to maximize $V(m_1, x_1)$, for doing this positively affects both her chances of being retained and the fee she gets in case she is retained.

Before we move on to the equilibrium analysis, it is worth commenting on the specific features of our setting, which combines a binary action with the reputational mechanism

[4]Note that w_1 plays no role and could be set equal to zero, while w_2 is instrumental to generate reputational concerns that, conditional on the expert being retained, are continuously increasing in the levels of reputations $\widehat{\alpha}$ and $\widehat{\gamma}$.

[5]We make this simplifying assumption for the sake of exposition. Allowing the expert to receive only a share of the expected value of her information does not affect the results.

described above. First, in terms of sorting, this structure allows the decision maker to fully exploit what she learns about the ability and the integrity of the incumbent at the end of period $t = 1$. Moreover, because our main focus is on welfare, adopting this structure significantly reduces the computational complexity with respect to a model with continuous actions.[6]

2.2 Equilibrium Analysis: Preliminaries

We use the concept of perfect Bayesian equilibrium and focus on informative equilibria defined as equilibria in which, in each period, the decision maker learns something decision-relevant from the expert's messages.

In this section, we provide a descriptive characterization of these equilibria, a formal analysis of which is relegated to the Appendix. The first thing to observe is that in any informative equilibrium, in each period, the expert's message must reveal some information about the state of the world.[7] This implies that in any informative equilibrium, m_t makes DM change her belief about x_t.[8] Because in our setting $R_t(1,1) = -R_t(1,0)$ and $\Pr(x_t = 1) = \frac{1}{2}$, it is then true that in any informative equilibrium, DM chooses $a_t(m_t) = m_t$.[9] With this in mind, we proceed by backward induction.

[6]In terms of sorting, this structure makes the model qualitatively equivalent to the model with continuous action and quadratic loss function used, for example, by Sobel (1985) and Morris (2001). In particular, in both those settings, DM takes an action based on the expected correctness of the expert's information, which depends on the expert's updated reputation. However, while in the continuous action model, sorting involves choosing a continuous action that minimizes expected loss, in our setting, it involves replacing an incumbent.

[7]This result is implied by Lemma 5(i) in the Appendix. Intuitively, because in the second period, learning about ability or integrity is no longer decision relevant for the future, any informative equilibrium must involve DM learning something about x_2. In the proof of Lemma 5(i) we further show that any equilibrium strategy profile in which the expert does not reveal any information about x_1 must necessarily be a "babbling" strategy, also implying that no learning occurs about either ability or integrity.

[8]Without loss of generality, we restrict attention to informative equilibria in which DM interprets message 1 to be (weakly) correlated with signal 1 and hence state 1.

[9]Put differently, in this model if an equilibrium is informative, it is also persuasive. With discrete actions and a prior that is not fair, an informative equilibrium may not be persuasive. For example, if either the prior on the state is extreme or the return in one state is extreme, a message by the expert may induce DM to revise her beliefs about the state. However, this revision may not be sufficient to induce DM to choose the action recommended by the expert.

2.2.1 The Second Period

Reporting strategies and DM's action. An expert that is active in the second and last period is not concerned about her reputation. For an unbiased expert with no preferences in favor of a particular action, any strategy is a continuation equilibrium. In line with the rest of the literature on career concerns, we focus on the continuation equilibrium in which an unbiased expert acts in the interest of DM and thus truthfully reveals her signal.[10] In this equilibrium, messages contain some information about the state of the world. Hence, DM chooses $a_2(m_2) = m_2$, and a biased expert reports $m_2 = 1$ regardless of her signal to induce action $a_2 = 1$.

The value function. Having pinned down the reporting strategies of biased and unbiased experts in the second period, we can now easily derive the value function $V(m_1, x_1)$, which represents the value of the incumbent's information in the second period. Note that this value is equal the payoff that DM expects to attain in the second period thanks to the information of the incumbent. Since at the moment of calculating this expected payoff, DM knows the incumbent's reputations, we have that:

$$V(m_1, x_1) \equiv E\left[R_2(a_2, x_2) \mid \widehat{\gamma}(m_1, x_1), \widehat{\alpha}(m_1, x_1)\right].$$

Given the equilibrium strategies that biased and unbiased experts use in the second period, it then follows that:

$$V(m_1, x_1) = \frac{r}{2}\widehat{\gamma}(m_1, x_1)\left[2\widehat{q}(m_1, x_1) - 1\right]. \tag{1}$$

It is straightforward to verify that $V(m_1, x_1)$ is strictly increasing in the incumbent's reputations $\widehat{\gamma}(m_1, x_1)$ and $\widehat{\alpha}(m_1, x_1)$.

DM's retaining strategy. At the beginning of the second period, DM chooses whether to retain the incumbent or hire a new expert. Again, using the second-period equilibrium strategies outlined at the beginning of this section, we obtain that the value of the

[10] Note that this equilibrium is the most informative in the Blackwell sense, but it is not unique. Indeed, any strategy profile that involves the unbiased expert revealing her signal with probability between 0 and 1 gives rise to an informative equilibrium that is obviously less informative than the one in which the unbiased expert truthfully reveals her signal. As our analysis focuses on first-period behavior, selecting this most informative continuation equilibrium is without loss of generality.

information of a new expert reads:

$$V \equiv E\left[R_2(a_2, x_2)\right] = \frac{r}{2}\gamma\left(2q - 1\right). \tag{2}$$

Given the analysis above, it should be apparent that at the beginning of the second period, DM retains the incumbent expert whenever $V(m_1, x_1) \geq V$ and replaces her with a new expert otherwise.[11]

2.2.2 The First Period

We are now ready to analyze the reporting strategies of biased and unbiased experts in the first period. In doing so, we assume that the continuation equilibrium described above is played.

First, let us define function $\imath(m_1, x_1)$ as follows:

$$\imath(m_1, x_1) = \begin{cases} 1 \text{ if } V(m_1, x_1) \geq V \\ 0 \text{ otherwise.} \end{cases} \tag{3}$$

Then, for a biased expert with signal s_1, the expected payoff of choosing message m_1 reads:

$$(1 - \delta_E)\left[w_1 + a(m_1)\right] + \delta_E \sum\nolimits_{x_1} \Pr(x_1 \mid s_1)\left[V(m_1, x_1) + 1\right]\imath(m_1, x_1). \tag{4}$$

For an unbiased expert with signal s_1, the expected payoff of choosing message m_1 reads:

$$(1 - \delta_E)\, w_1 + \delta_E \sum\nolimits_{x_1} \Pr(x_1 \mid s_1)\left[V(m_1, x_1)\right]\imath(m_1, x_1). \tag{5}$$

Biased and unbiased experts will respectively choose m_1 to maximize expressions (4) and (5). It is worth noticing that while m_1 affects both the current and the future payoff of a biased expert, it only affects the future payoff of an unbiased expert. In other words, while a biased expert has both current and reputational incentives, an unbiased expert only has reputational concerns.

[11] Because both q and $\widehat{q}(m_1, x_1)$ are greater than $\frac{1}{2}$ (i.e., in expectation the expert is better informed than DM), both $\widehat{V}(m_1, x_1)$ and V are strictly positive. Thus, DM always finds it optimal to consult an expert in period 2.

It turns out that a multitude of informative first-period reporting strategies is consistent with the continuation equilibrium outlined in the previous subsection. In what follows, we use the expression *truthtelling equilibrium* (or simply *truthtelling*) to denote an equilibrium in which both biased and unbiased experts truthfully reveal their signals in the first period; we instead use the expression *informative misreporting equilibrium* (or simply *misreporting equilibrium*) to denote an equilibrium in which experts' signals are partially disclosed in the first period.

As we are interested in analyzing whether misreporting equilibria have the potential to increase welfare with respect to truthtelling, it is convenient to introduce the basic tools of the welfare analysis at this stage, and then proceed with the characterization of the various informative equilibria. In doing so, we are implicitly assuming that both truthtelling and informative misreporting equilibria exist. We indeed show that this is true in Sections 4 and 5.

3 Welfare: Discipline versus Sorting

As mentioned in Section 2, we focus on the decision maker's welfare and thus on the ex-ante expected payoff of DM in a given equilibrium σ, defined as follows:[12]

$$E_0^{\sigma}\left[\Pi_{DM}\right] = \left(1-\delta_{DM}\right) E_0^{\sigma}\left[R_1(a_1,x_1)\right] + \delta_{DM} E_0^{\sigma}\left[R_2(a_2,x_2)\right]. \tag{6}$$

As a first step towards analyzing welfare, it is useful to identify two distinct effects that emerge in equilibrium, namely the *discipline* and *sorting* effects. The discipline effect arises in the first period, when reputational concerns induce an expert to reveal some of her information about the state of the world. The sorting effect arises at the end of the first period, when DM learns something about the incumbent's ability and integrity after observing m_1 and x_1. While the discipline effect positively affects the expected payoff of the first period decision (i.e., $E_0^{\sigma}\left[R_1(a_1,x_1)\right]$), the sorting effect positively affects the expected payoff of the second period decision (i.e., $E_0^{\sigma}\left[R_2(a_2,x_2)\right]$). A truthtelling equilibrium always involves greater discipline and thus a higher expected utility of current decisions than any

[12]Throughout the paper, equilibrium values in a particular equilibrium will be denoted with a superscript representing the name of that particular equilibrium.

other misreporting equilibrium. However, when we compare a misreporting equilibrium with a truthtelling equilibrium in terms of how much the decision maker learns about the integrity and the ability of the incumbent expert, results are not so straightforward.

Let ME denote an informative misreporting equilibrium and TT denote a truthtelling equilibrium. We then say that a given equilibrium ME improves sorting with respect to TT if and only if $E_0^{ME}\left[R_2(a_2,x_2)\right] > E_0^{TT}\left[R_2(a_2,x_2)\right]$. To analyze under what conditions this inequality is satisfied, we formally define distinct measures of learning for both integrity and ability.

Definition 1 *$|\widehat{\gamma}^{\sigma}(m_1,x_1)-\gamma|$ and $|\widehat{\alpha}^{\sigma}(m_1,x_1)-\alpha|$ respectively measure the amount of learning about integrity and ability in a putative equilibrium σ when realization (m_1,x_1) is observed.*

The following proposition then establishes a general property of informative misreporting equilibria that suggests that lies may have a positive effect.

Proposition 1 *For every (m_1,x_1),*

$$\left|\widehat{\gamma}^{ME}(m_1,x_1)-\gamma\right| > \left|\widehat{\gamma}^{TT}(m_1,x_1)-\gamma\right| = 0,$$

and

$$\left|\widehat{\alpha}^{ME}(m_1,x_1)-\alpha\right| \leq \left|\widehat{\alpha}^{TT}(m_1,x_1)-\alpha\right|,$$

with strict inequality for at least one (m_1,x_1) {Proof in the Appendix}.

In other words, proposition 1 suggests that relative to truthtelling, all informative misreporting equilibria lead to more learning about integrity and less learning about ability. To see this, note that a biased expert has the same probability that an unbiased expert has of receiving any given signal. Because biased and unbiased experts use the same reporting strategy in a truthtelling equilibrium, any given message is as likely to come from one type or the other. Therefore, messages are completely uninformative about integrity. On the contrary, informative misreporting equilibria are characterized by biased and unbiased experts using different reporting strategies. This implies that, in equilibrium, each message is sent more frequently by one type of expert or the other. Hence, the message in itself allows DM to learn something about the expert's integrity. For example, if a biased doctor recommends surgery more often than an unbiased doctor, receiving a surgery recommendation will rationally lead the patient to believe that the doctor is more likely to be biased than when she is prescribed a more conservative treatment. As for ability, note that in a truthtelling equilibrium, observing the expert's recommendation is equivalent to observing her information. Hence, the decision maker is in the best position to evaluate the quality of the expert's signals. This is not the case in a misreporting equilibrium. Indeed, as there is some lying, messages do not fully reflect the information of the expert. Hence, inference about the ability of the expert is less sharp than in TT for at least some realizations (m_1, x_1).

Having established that informative misreporting equilibria lead to more learning about preferences does not imply that these equilibria will necessarily lead to better expected decisions in the future (i.e., to better sorting) than truthtelling. Considering the expression for $V(m_1, x_1)$ given by (1), the value of the expert's information in the second period depends on both ability and integrity. To clarify, and continuing with the patient-doctor example, if a patient learns that a doctor is unbiased without learning enough about the doctor's ability, it is not obvious that the patient will receive more informed medical advice in the future.

In the following sections, we show that there are several cases in which informative misreporting equilibria actually lead to better sorting than truthtelling. Considering the expression for $E_0^{\sigma}[\Pi_{DM}]$ given by (6), it then becomes clear that a misreporting equilibrium with better sorting has the potential to dominate truthtelling. Whether this occurs or not depends on DM's preferences for the future versus the present as established in the following lemma.

Lemma 1 *For any informative misreporting equilibrium that improves sorting with respect to truthtelling, there always exists a $\delta^*_{DM} \in (0,1)$, such that the misreporting equilibrium increases (decreases) $DM's$ ex-ante expected utility with respect to truthtelling if $\delta_{DM} > \delta^*_{DM}$ ($\delta_{DM} < \delta^*_{DM}$).*

Proof. For any putative informative misreporting equilibrium ME and truthtelling equilibrium TT, $E_0^{ME}[R_1(a_1, x_1)] < E_0^{TT}[R_1(a_1, x_1)]$. If ME improves sorting, we have that

$E_0^{ME}[R_2(a_2, x_2)] > E_0^{TT}[R_2(a_2, x_2)]$. As $E_0^{\sigma}(\Pi_{DM})$ is monotonic in δ_{DM}, this completes the proof. ■

We are now ready to complete the analysis of Section 2 and characterize the informative equilibria of our game. This amounts to characterizing the first-period reporting strategies of biased and unbiased experts. For this reason, in what follows, when we analyze the behavior of biased and unbiased experts, we refer to first-period behavior.

For the sake of exposition, we divide informative equilibria into two main classes: *i)* equilibria in which the unbiased expert truthfully reports her signals in the first period; and *ii)* equilibria in which the unbiased expert misreports. In Section 4, we begin by analyzing the first class of equilibria. We then focus on the second class of equilibria in Section 5. For each misreporting equilibrium that we identify, we compare how it fares in terms of welfare with respect to truthtelling.

4 Truthful Reporting by the Unbiased Expert

In this section, we consider informative equilibria in which the unbiased type reports truthfully. This allows us to show the existence of truthtelling equilibria and then illustrate the main results of the paper. The following proposition characterizes the class of informative equilibria in which the unbiased expert truthfully reports her signals, by dividing them into two subclasses which we label TT and MB.

Proposition 2 *Each equilibrium in which U truthfully reports her signals belongs to one of the following two subclasses:*

i) Truthtelling (TT), in which B truthfully reveals her signals;

ii) Misreporting Biased (MB), in which B reports signal $s_1 = 1$ truthfully, and signal $s_1 = 0$ with probability $0 < \lambda_{B,0} < 1$.

In both subclasses, DM retains the incumbent expert if and only if $m_1 = x_1$.

{*Proof in the Appendix*}.

It is worth noticing that a TT equilibrium could never be supported if reputational concerns were only related to preferences, as in Morris (2001). It is the presence of a second dimension of reputation (i.e. reputation for ability) that creates the right incentives to fully reveal information about the state of the world.[13]

To gather a better understanding of how each of these equilibria arise, consider Figure 1, which shows how DM's beliefs $\widehat{\alpha}(m_1, x_1)$ and $\widehat{\gamma}(m_1, x_1)$ vary as a function of $1 - \lambda_{B,0}$ (i.e., the probability with which the biased expert misreports $s_1 = 0$). The case $\lambda_{B,0} = 1$ identifies a TT equilibrium, while the case $0 < \lambda_{B,0} < 1$ identifies an MB equilibrium. Note that when $\lambda_{B,0} = 1$, we have that:

$$\widehat{\gamma}^{TT}(m_1, x_1) = \gamma \text{ for all } (m_1, x_1),$$

$$\underline{\alpha} \equiv \widehat{\alpha}^{TT}(1, 0) = \widehat{\alpha}^{TT}(0, 1) < \alpha < \widehat{\alpha}^{TT}(1, 1) = \widehat{\alpha}^{TT}(0, 0) \equiv \overline{\alpha}.$$

Hence, as shown in proposition 1, messages have no impact on the reputation for being unbiased. However, in a TT equilibrium, reporting a correct (incorrect) message causes the expert to establish the highest (lowest) reputation for being smart. To ease notation, throughout the paper, we will use $\overline{q}$ to denote the value of $\widehat{q}(m_1, x_1)$ corresponding to $\overline{\alpha}$.

Given the above values of the reputations, we may verify that in a truthtelling equilibrium the following holds true:

$$\underline{V} \equiv V^{TT}(1, 0) = V^{TT}(0, 1) < V < V^{TT}(1, 1) = V^{TT}(0, 0) \equiv \overline{V}. \tag{7}$$

Relation (7) implies that in a TT equilibrium, DM retains the incumbent if she makes a correct recommendation and replaces her if she makes a mistake. Because signals are on average informative, truthfully reporting a signal maximizes the chances of providing a correct recommendation and hence being retained. For this reason, for an unbiased expert who is solely concerned about the impact of m_1 on her continuation payoff, always reporting $m_1 = s_1$ is consistent with the equilibrium. The same incentive applies to a biased expert when δ_E is sufficiently large, in other words when the biased expert is more concerned about the continuation payoff than the current payoff. In the appendix, we show that there always exists a scalar $\underline{\delta}_E^{TT} \in (0, 1)$ such that if $\delta_E \geq \underline{\delta}_E^{TT}$, then a biased expert always reports truthfully. Thus, a TT equilibrium exists if and only if a biased expert is sufficiently concerned about her career prospects. All this is driven by the reputational

[13] Section 7 explores this issue in further detail.

concern for ability, since in a truthtelling equilibrium there is no variation in the reputation for integrity.

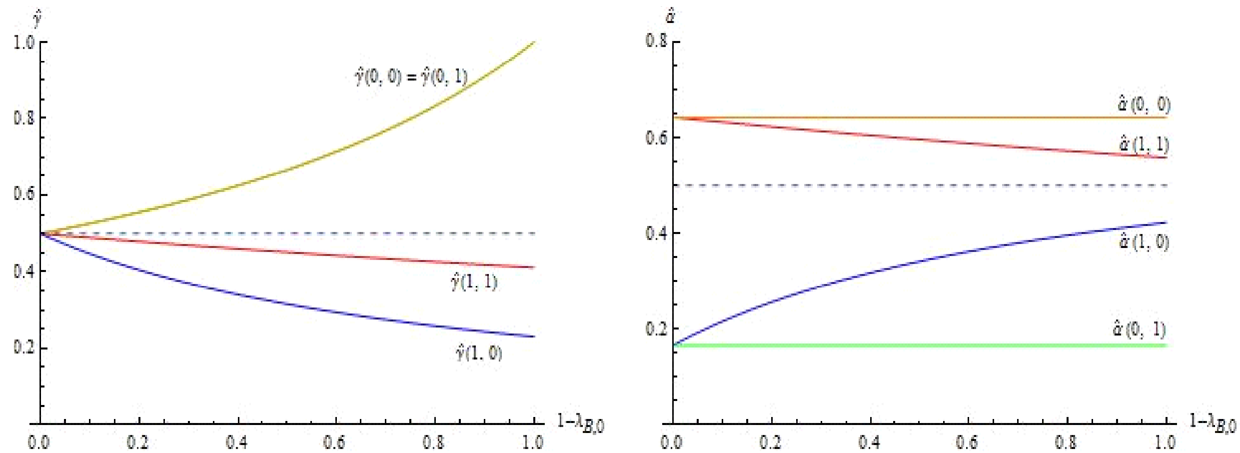

Figure 1: Reputation values as a function of B's probability of misreporting $(1-\lambda_{B,0})$

When $\delta_E < \underline{\delta}_E^{TT}$, the career concerns of a biased expert are not sufficiently high to induce her to truthfully report all her signals. In particular, a biased expert will be tempted to lie when receiving $s_1 = 0$. In the appendix, we show that there always exists a scalar $\underline{\delta}_E^{MB} \in \left(0, \underline{\delta}_E^{TT}\right)$ such that if $\underline{\delta}_E^{MB} \leq \delta_E < \underline{\delta}_E^{TT}$, the MB equilibria described in proposition 2 exist. In these equilibria, while an unbiased expert reports truthfully, a biased expert truthfully reports $s_1 = 1$ and partially reveals $s_1 = 0$.

We can again turn to Figure 1 for an intuition of how MB equilibria arise. From figure 1, it is apparent that when $0 < \lambda_{B,0} < 1$, the decision maker's beliefs $\widehat{\alpha}(m_1, x_1)$ and $\widehat{\gamma}(m_1, x_1)$ satisfy the following relations:

$$\widehat{\gamma}^{MB}(1,0) < \widehat{\gamma}^{MB}(1,1) < \gamma < \widehat{\gamma}^{MB}(0,1) = \widehat{\gamma}(0,0)^{MB},$$

$$\underline{\alpha} = \widehat{\alpha}^{MB}(0,1) < \widehat{\alpha}^{MB}(1,0) < \alpha < \widehat{\alpha}^{MB}(1,1) < \widehat{\alpha}^{MB}(0,0) = \overline{\alpha}.$$

In other words, in an MB equilibrium, observing $m_1 = 0$ ($m_1 = 1$) increases (reduces) the value of $\widehat{\gamma}(m_1, x_1)$ above (below) the prior γ. This occurs because in MB, an unbiased expert reports $m_1 = 0$ more often than a biased expert, and $m_1 = 0$ ($m_1 = 1$) thus conveys information about the expert's likelihood of being unbiased (biased). Concerning ability, because messages partially reflect the private signal of the expert, the comparison of m_1 with x_1 allows DM to infer something about the ability of the expert. Note that since there is some misreporting there are some realizations (m_1, x_1) after which this inference is less sharp than in a TT equilibrium. However, as in TT, an ex-post correct (wrong) message increases (decreases) the value of $\widehat{\alpha}(m_1, x_1)$ above (below) the prior α, thereby providing an incentive to report truthfully. Following this argument, it should be intuitive that in an MB equilibrium, the expected reputational reward from reporting $m_1 = 0$ when $s_1 = 0$ is observed is relatively large since it enhances both reputations. Indeed, it is this substantial reputational reward that eventually offsets the low value of δ_E and preserves the incentive of a biased expert to partially reveal $s_t = 0$ instead of disregarding it completely.

4.1 When Can Misreporting Be Preferred to Truthtelling?

We now compare how MB equilibria fare in terms of welfare with respect to TT equilibria. Figure 1 shows us the pattern of learning about ability and integrity in the two cases. Not surprisingly, this is consistent with the result of proposition 1. MB equilibria (i.e., equilibria with misreporting) lead to more learning about integrity and less learning about ability with respect to truthtelling. We know by Lemma 1 that if this learning pattern leads to better sorting, MB has the potential to improve welfare with respect to TT. The following proposition highlights how MB equilibria fare with respect to TT ones in terms of sorting.

Proposition 3 *There always exists a scalar $\alpha^* \in (0, 1)$, such that MB improves sorting with respect to TT if and only if $\alpha > \alpha^*$.{Proof in the Appendix}.*

For an intuition of this result recall that, as shown in proposition 1, TT guarantees the sharpest learning about ability yet no learning about integrity, while MB allows for some learning about both ability and integrity. Since when α is high there is little scope for learning about ability, this allows MB to improve sorting with respect to TT.

To see this, let $\Pr(m_1, x_1 \mid \sigma)$ denote the ex-ante probability that realization (m_1, x_1) is observed given that equilibrium σ is played. Then, consider the following expressions representing the ex-ante second-period expected payoffs in MB and TT:

$$\begin{aligned} E_0^{MB}\left[R_2(a_2, x_2)\right] &= \Pr(1,1|MB)V^{MB}(1,1) + \Pr(0,0|MB)V^{MB}(0,0) \\ &\quad + \Pr(0,1|MB)V + \Pr(0,1|MB)V, \end{aligned} \tag{8}$$

$$\begin{aligned} E_0^{TT}\left[R_2(a_2, x_2)\right] &= \Pr(1,1|TT)\overline{V} + \Pr(0,0|TT)\overline{V} \\ &\quad + \Pr(1,0|TT)V + \Pr(0,1|TT)V. \end{aligned} \tag{9}$$

Proposition 3 states that when α is sufficiently high, $E_0^{MB}\left[R_2(a_2, x_2)\right] - E_0^{TT}\left[R_2(a_2, x_2)\right] > 0$. Note that this difference can be decomposed into two components. First, consider the difference between the bites of (8) and (9) that refer to the events in which the expert makes a mistake and hence is fired (i.e., events in which $m_1 \neq x_1$). We denote this value as the *replacement component*, which can be written as follows:

$$\begin{aligned} &\Pr(1,0|MB)V + \Pr(0,1|MB)V - \left[\Pr(1,0|TT)V + \Pr(0,1|TT)V\right] = \\ = \; &\left[\Pr(m_1 \neq x_1, B|MB) - \Pr(m_1 \neq x_1, B|TT)\right] V > 0. \end{aligned} \tag{10}$$

Expression (10) highlights that the replacement component is positive. This occurs because the probability of replacing an unbiased expert is the same in both equilibria, while the probability of correctly replacing a biased expert is strictly higher in MB than in TT. Indeed, since in MB the biased expert misreports with positive probability, her chances of making a mistake are larger than in TT.

Now, consider the difference between the bites of (8) and (9) that refer to the events in which the expert provides a correct recommendation (i.e., events in which $m_1 = x_1$). We denote this value as the *continuation component,* which reads:

$$\begin{aligned} &\Pr(1,1|MB)V^{MB}(1,1) + \Pr(0,0|MB)V^{MB}(0,0) + \\ &- \left[\Pr(1,1|TT)\overline{V} + \Pr(0,0|TT)\overline{V}\right]. \end{aligned}$$

After replacing the equilibrium values of $V^{MB}(m_1, x_1)$ and $\overline{V}$ and simplifying, the previous expression becomes:

$$\frac{r}{2}q\gamma\left[\widehat{q}^{MB}(1,1) - \overline{q}\right] < 0. \tag{11}$$

Since TT is the equilibrium with the sharpest learning about ability, $\overline{q} \geq \widehat{q}^{MB}(1,1)$ implying that expression (11) is always negative. This means that, for MB to exhibit a stronger sorting effect with respect to TT, (10) must be large enough compared to (11). When α is sufficiently high, there is little scope for learning about ability. Consequently, the difference $\widehat{q}^{MB}(1,1) - \overline{q}$ becomes small and so does expression (11). At the same time, expression (10) remains strictly positive since $\Pr(m_1 \neq x_1, B|MB)$ is strictly larger than $\Pr(m_1 \neq x_1, B|TT)$ due to misreporting by the biased expert in MB.[14] This result does not symmetrically apply when α is sufficiently close to zero because in this case both (10) and (11) tend to zero. For (11), this occurs for the same reasons we mentioned when α is high. For (10), notice that for small values of α, signals tend to be uninformative. Therefore, the probability of correctly replacing a biased expert that makes an incorrect evaluation tends to $1/2\Pr(B)$ in both MB and TT. This implies that (10) tends to zero.

Proposition 3 suggests that it may not always be the case that TT is the welfare maximizing equilibrium. While TT allows for a higher expected utility of current decisions (discipline effect), MB may imply better expected decisions in the future thanks to a stronger sorting effect. As mentioned in Lemma 1, if DM is sufficiently concerned about future decisions, then MB may indeed improve welfare with respect to TT.

5 Misreporting by the Unbiased Expert

So far we have restricted our analysis to the class of equilibria in which an unbiased expert truthfully reports all her signals. However, there also exist equilibria in which the unbiased expert misreports. These equilibria have the flavor of the political correctness equilibria described by Morris (2001), since the unbiased expert lies and sends a specific message more often than the biased expert to signal her type to the decision maker. The following proposition characterizes the class of informative equilibria in which the unbiased expert misreports, by dividing them into two subclasses which we label MU and TMU.

Proposition 4 *Each equilibrium in which U misreports belongs to one of the following two*

[14]More formally, in the Appendix we show that there always exists an α strictly less than 1 for which both MB and TT exist, and above which the replacement component always dominates the continuation component. Note also that this result holds true for every $\gamma \in (0,1)$. Therefore, the result of Proposition 3 does not rely on uncertainty about integrity being greater than uncertainty about ability.

subclasses:

i) Misreporting Unbiased (MU), in which U partially reveals one signal and truthfully reports the other signal, and DM retains the incumbent if and only if $m_1 = x_1$;

ii) Total Misreporting Unbiased (TMU), in which U always sends $m_1 = 0$ regardless of her signal, and DM retains the incumbent if $m_1 = x_1 = 0$, and replaces her if $m_1 = 1$.

In both subclasses, B's strategy must be such that the message that is falsely reported by U is sent more often by U than by B.

{*Proof in the Appendix*}.

Notice that misreporting a signal implies sending a message that is more likely to be incorrect ex-post and hence reduces the expert's expected reputation for ability. Thus, the only reason for U to lie is that lying brings about a sufficient increase in the reputation for integrity. This occurs if, in equilibrium, the message that is falsely reported by U is sent more often by U than by B, so that such a message "signals" that the sender is more likely to be unbiased. This is exactly what happens in the MU and TMU equilibria of proposition 4.

Note that MU equilibria include two cases, each characterized by a different misreporting behavior by U. In the first case, U partially reveals $s_1 = 1$ and truthfully reveals $s_1 = 0$; in the second case, the opposite holds.[15] TMU equilibria instead represent the more extreme cases in which the unbiased expert never communicates the evaluation favored by the biased expert. These equilibria are supported by a very conservative strategy of the decision maker. Indeed, DM ignores the ex-post correctness of a message and retains the incumbent if $m_1 = 0$ is ex-post correct and replaces her whenever $m_1 = 1$ is observed.[16]

A natural question is whether equilibria in which an unbiased expert lies have the potential to improve sorting and hence the expected utility of DM with respect to truthtelling equilibria.

[15] Notice that MU equilibria in which U partially reveals $s_1 = 0$ and truthfully reveals $s_1 = 1$ are specific to our setting. Here, U lies by falsely reporting $m_1 = 1$. To support this equilibrium, B's strategy must be such that if she misreports $s_1 = 0$, she must do so with higher probability than U so that $m_1 = 1$ is eventually sent more often by U than by B. One may wonder how it can be that in equilibrium, B sends her favorite message less often than U. In our setting, this can occur because the bias is "relation specific". This implies that B benefits from DM choosing $a_2 = 1$ if and only if B has been retained by DM. As in these equilibria the expert is retained if and only if $m_1 = x_1$, B has some incentive to report $m_1 = 0$ after observing $s_1 = 0$ because doing so maximizes the probability that the message is ex-post correct.

[16] As we show in the Appendix, there exist TMU equilibria in which DM retains the expert after $(m_1, x_1) = (0, 1)$ as well as TMU equilibria in which DM replaces the expert afetr $(m_1, x_1) = (0, 1)$.

5.1 Equilibria in which an Unbiased Expert Lies that Are *Never* Preferred to Truthtelling

We begin by considering MU equilibria. In particular, we would expect those MU in which B reports truthfully to have the potential to improve sorting with respect to TT, since the structure of the reporting strategy is the same as in MB. Rather surprisingly, we find that this is not the case and that they are always dominated by truthtelling.[17] The following proposition highlights this result:

Proposition 5 *MU equilibria in which B truthfully reports can never improve sorting with respect to TT.* {*Proof in the Appendix*}.

While we confine the formal proof of proposition 5 to the appendix, we now provide an intuition of why U's misreporting negatively affects sorting. Intuitively, this occurs because in the attempt to signal her type, U misreports, and this gives rise to a learning pattern that reduces DM's chances of consulting an unbiased expert of high ability in the second period. To see this, it is convenient to compare the sorting effect that arises in MU equilibria characterized by U partially revealing $s_1 = 1$ and truthfully revealing $s_1 = 0$, with the sorting effect that arises in MB.[18] To make this comparison meaningful, let us assume that the expected amount of lying in the first period is the same in the two equilibria. This is the case when $\gamma = \frac{1}{2}$ and the probability with which the biased expert and the unbiased expert lie is the same (i.e., $\lambda_{U,1} = \lambda_{B,0} = \lambda$).[19] Under these assumptions, any difference between the two equilibria is not driven by the fact that one type of expert lies more or less than the other type.

First, notice that if the expected amount of lying in the first period is the same in both equilibria, the probability that $m_1 \neq x_1$ is also the same. Since the payoff associated with events in which $m_1 \neq x_1$ is also the same in both equilibria (i.e., V), the difference in sorting between MB and TT can be entirely explained by considering what happens

[17] With respect to MU equilibria in which B randomizes, note that as U's probability of misreporting tends to zero, these equilibria converge to MB and thus have the potential to improve sorting with respect to truthtelling.

[18] A similar argument applies when comparing MU equilibria in which U partially reveals $s_1 = 0$ and truthfully reveals $s_1 = 1$ with MB equilibria.

[19] The expected amount of lying is $(1 - \gamma)\lambda_{B,0}$ in MB and $\gamma\lambda_{U,1}$ in these MU equilibria.

when $m_1 = x_1$ (i.e., by the continuation component). Indeed, we have that:

$$\begin{aligned} & E_0^{MB}\left[R_2(a_2, x_2)\right] - E_0^{MU}\left[R_2(a_2, x_2)\right] = \\ = \; & \Pr(1,1|MB)V^{MB}(1,1) + \Pr(0,0|MB)V^{MB}(0,0) + \\ & - \Pr(1,1|MU)V^{MU}(1,1) - \Pr(0,0|MU)V^{MU}(0,0). \end{aligned}$$

Notice that in both equilibria, message $m_1 = 0$ signals that the expert is likely to be unbiased, and message $m_1 = 1$ signals that the expert is likely to be biased. Instead, learning about ability takes place in a different way in each of the two cases. In MB equilibria, reporting strategies are such that message $m_1 = 0$ perfectly reveals signal $s_1 = 0$, while message $m_1 = 1$ only imperfectly reveals signal $s_1 = 1$. This implies that reputation for ability increases more after realization $(m_1, x_1) = (0, 0)$ than after realization $(m_1, x_1) = (1, 1)$ (i.e., $\widehat{\alpha}^{MB}(0,0) > \widehat{\alpha}^{MB}(1,1)$). The opposite occurs in MU because message $m_1 = 1$ perfectly reveals signal $s_1 = 1$, and message $m_1 = 0$ only imperfectly reveals signal $s_1 = 0$ (i.e., $\widehat{\alpha}^{MU}(1,1) > \widehat{\alpha}^{MU}(0,0)$). It is easy to see that when the expected amount of lying is the same in the two equilibria, we have that:

$$\alpha < \widehat{\alpha}^{MU}(0,0) = \widehat{\alpha}^{MB}(1,1) < \widehat{\alpha}^{MU}(1,1) = \widehat{\alpha}^{MB}(0,0) = \overline{\alpha}.$$

This implies that:

$$q < \overline{q}_{low} \equiv \widehat{q}^{MU}(0,0) = \widehat{q}^{MB}(1,1) < \widehat{q}^{MU}(1,1) = \widehat{q}^{MB}(0,0) = \overline{q}.$$

Now, using equilibrium values and simplifying, we can rewrite the difference in sorting between MB and TT as follows:

$$E_0^{MB}\left[R_2(a_2, x_2)\right] - E_0^{MU}\left[R_2(a_2, x_2)\right] = \gamma(1-\lambda)q(2\overline{q}-1) - \gamma(1-\lambda)(1-q)(2\overline{q}_{low}-1) > 0. \quad (12)$$

The first term is positive since after realizations (m_1, x_1) that reveal more information on ability (i.e., $(0, 0)$ in MB; and $(1, 1)$ in MU), it is more likely that the expert is unbiased in MB than in MU. This is so because in both equilibria, $m_1 = 0$ is sent more often by U than by B, and thus $m_1 = 0$ ($m_1 = 1$) signals that the expert is likely to be unbiased (biased). More specifically, in the first case, the unbiased expert always truthfully reports her signal in the state that is more informative on ability, while in the second case, this occurs with probability $\lambda < 1$. The second term is negative due to the fact that realizations (m_1, x_1) that reveal less information on ability (i.e., $(1,1)$ in MB; and $(0,0)$ in MU) are more frequently associated with an unbiased expert in MU, since in this equilibrium, when receiving signal $s_1 = 1$, U misreports with probability $(1-\lambda)$ by sending message 0, whereas U never misreports in MB.

Overall, the positive term always outweighs the negative one so that (12) is greater than zero. To see this, first recall that the positive term is associated to more learning about ability, in other words, as mentioned previously $\overline{q} > \overline{q}_{low}$. In addition, the negative term is assigned a smaller probability with respect to the positive term. Indeed, since the negative term represents cases in which there is misreporting, the evaluation is less likely to be correct and the chances of being retained are equal to $(1-q)$, while for the positive term, as it represents situations in which the expert reports truthfully, , her odds of being retained are equal to q.

5.2 Equilibria in which an Unbiased Expert Lies that *May* Be Preferred to Truthtelling

We now consider the more conservative equilibria in which DM never retains an expert that reports $m_1 = 1$. It turns out that these equilibria may improve sorting relative to truthtelling equilibria. The following proposition states this result.

Proposition 6 *There exists a non-empty set of values of p and α for which it is always possible to find a scalar $\gamma^* \in (0,1)$ such that for any $\gamma < \gamma^*$, there exists a TMU equilibrium that improves sorting with respect to TT. {Proof in the Appendix}.*

To identify equilibria that satisfy proposition 6, we consider the TMU equilibria in which DM retains the incumbent if and only if $(m_1, x_1) = (0,0)$ is observed, and B always truthfully reports her signals.[20] As we show in the appendix, these particular TMU equilibria may improve sorting with respect to TT when DM faces higher odds of encountering a

[20]We focus on these particular equilibria only for the sake of exposition. Indeed, also for the TMU equilibria in which the expert is retained even after $(m_1, x_1) = (0,1)$, it is possible to show that there exists a non-empty space of parameters for which these can improve sorting with respect to TT.

biased expert, that is, when γ is relatively low. *To understand why, notice that as γ tends to 1, we have that in both equilibria DM is likely to end up retaining an expert that is unbiased. However, since TT allows for sharper learning about ability, the unbiased expert is more likely to be smart in TT than in TMU. This makes TT superior to TMU when γ tends to 1.* On the contrary, when γ is below a certain threshold, the conservative replacement strategy implied by this particular TMU equilibrium allows DM to better discriminate biased versus unbiased experts while also learning something about ability. To see this, we can break up the net welfare gain of TMU with respect to TT into the usual two components, namely the bite that refers to the events in which the expert makes a mistake and is fired (i.e., the replacement component):

$$\begin{aligned} & [\Pr((m_1, x_1) \neq (0,0))|TMU) - \Pr(m_1 \neq x_1|TT)]\, V = \\ = \; & \frac{r}{2}\gamma \left[(2q-1)\, \frac{1}{2} \left((1-\gamma)q + \gamma\, (2q-1) \right) \right], \end{aligned} \tag{13}$$

and the bite that refers to the events in which the expert provides a correct recommendation and is retained (i.e., the continuation component):

$$\Pr(0,0|TMU)V^{TMU}(0,0) - \Pr(m_1 = x_1|TT)\overline{V} = \frac{r}{2}\gamma \left[\frac{1}{2} \left(2\widehat{q}_{00}^{TMU} - 1 \right) - q(2\overline{q} - 1) \right]. \tag{14}$$

Note that the replacement component (13) is always positive. This is because the probability of observing realizations (m_1, x_1) after which the incumbent expert is replaced is larger in this particular TMU than in TT. The decision maker is therefore more likely to fire the incumbent expert in the former rather than in the latter equilibrium. This is simply due to the fact that in this particular TMU: i) biased experts behave in the same way as in TT, but they are also replaced after realization $(1, 1)$; and ii) unbiased experts always send $m_1 = 0$ regardless of their signals, which makes the probability of making a mistake and hence being fired larger than in TT. Terms $(1-\gamma)q$ and $\gamma\,(2q-1)$ respectively capture the increase in the probability of firing a biased expert and the increase in the probability of firing an unbiased expert relative to TT.

Instead, the continuation component (14) is always negative. This is so for two reasons. First, in any equilibrium with misreporting, there is less learning about ability. Therefore, whenever DM retains an expert that made a correct evaluation, she is less certain about the expert's ability. This is captured by $\widehat{q}_{00}^{TMU} < \overline{q}$. Second, in TMU, experts are retained only when they provide a correct evaluation after sending message $m_1 = 0$, which occurs with probability $1/2$. On the other hand, in TT, experts are retained as long as they provide a correct evaluation independently of the message they send, which occurs with probability $q > \frac{1}{2}$.

Now notice that as γ decreases, the negative term $\left[\frac{1}{2}(2\widehat{q}_{00}^{TMU} - 1) - q(2\overline{q} - 1)\right]$ in expression (14) shrinks. This is because biased experts report truthfully in both equilibria, while unbiased experts do so only in TT. Therefore, as γ decreases, the fraction of experts that lie (tell the truth) in TMU decreases (increases) allowing DM to learn more about ability. All this is captured by q_{00}^{TMU} increasing and approaching $\overline{q}$ as γ decreases. Intuitively, as γ decreases, the expected quality of the expert that is retained tends to be the same in both equilibria. This reduces the advantage of TT over TMU. At the same time, as γ decreases, the positive term $\left[(2q-1)\,\frac{1}{2}\left((1-\gamma)q + \gamma(2q-1)\right)\right]$ in expression (13) increases. This occurs because in TT, the probability of firing an expert does not depend on γ, while in TMU, this probability is decreasing in γ since the chances of firing a biased expert are greater than those of firing an unbiased expert.[21] As we show in the proof, below a certain threshold of γ, the net advantage of TT over TMU in terms of a sharper learning about ability becomes small and is offset by the net advantage of TMU over TT in terms of better sorting out biased versus unbiased experts.

6 A Complete Mapping of Equilibria and Welfare Implications

So far, we have established that TT equilibria may sometimes be dominated by other equilibria that involve some degree of misreporting. To provide a more complete picture of our results, it is useful to present a mapping of all the equilibria based on the priors that represent the information environment. In particular, we characterize the equilibria with respect to the career concerns of the experts represented by parameter δ_E. This allows us to establish which types of equilibria may exist for the different regions of δ_E, to better

[21] In TMU, an unbiased expert only sends message $m_1 = 0$, while a biased expert sends both message $m_1 = 0$ and message $m_1 = 1$. Hence, it is more likely that realizations (m_1, x_1) after which an expert is fired arise from a biased expert rather than from an unbiased one.

comprehend in which cases truthful equilibria may or may not be welfare maximizing.

The first thing to notice is that informative equilibria exist whenever the expert assigns a high enough weight to future payoffs. Moreover, for all the values of δ_E for which informative equilibria exist, there is always a potential for multiplicity. Another relevant feature is that truthtelling equilibria may never coexist with MB, as shown in Section 2. The following proposition formally represents this situation:

Proposition 7 *There exist $\underline{\delta}_E, \overline{\delta}_E \in (0,1)$ with $\underline{\delta}_E < \overline{\delta}_E$ such that:*

a) For $\delta_E < \underline{\delta}_E$, no informative equilibria exist;

b) For $\delta_E \in (\underline{\delta}_E, \overline{\delta}_E)$, there always exists a non-empty set of informative equilibria that includes MB and at most both of the following: MU and TMU;

c) For $\delta_E \in (\overline{\delta}_E, 1)$, there always exists a non-empty set of informative equilibria that includes TT and at most both of the following: MU, and TMU.

{*Proof in the Appendix*}.

Although equilibrium multiplicity does not allow us to uniquely establish which equilibrium will be played, the welfare maximizing equilibrium represents the best possible outcome attainable for a given range of values of δ_E. This complete mapping of the equilibria allows us to state that truthtelling may not necessarily be welfare maximizing. In particular, the following general welfare results apply. First, when δ_E is sufficiently high, TT is feasible, and welfare maximizing as long as TMU does not exist. However, whenever TMU exists, it may dominate TT, as we observed in the previous section. Furthermore, when experts do not care enough about future payoffs, truthtelling breaks down, but there always exist other equilibria that involve some degree of misreporting (either MB or TMU) that may even generate higher levels of welfare with respect to truthtelling.

As a final observation, we consider the role that commitment may play in this setting. If we assume that DM can commit to a replacement strategy ex-ante and that non-babbling equilibria will always be played if they exist, then this device can, in some cases, function as a mechanism for eliminating welfare dominated equilibria. For instance, consider the case in which $\delta_E > \overline{\delta}_E$ and TMU, MU and TT equilibria all exist. By committing to one of the two strategies that are consistent with equilibrium, that is, either retaining the expert if $m_1 = 0$ and always replacing her if $m_1 = 1$, or retaining her if and only if $m_1 = x_1$, DM can respectively induce either TMU or one between MU and TT. In this setting, DM

should then adopt the first strategy whenever all TMU equilibria dominate the other two equilibria and the second strategy whenever the welfare-inferior equilibrium between MU and TT dominates all TMU equilibria. However, if for example some MU is dominated by some TMU when TT dominates all TMU, then committing to a replacement strategy may not be useful for eliminating dominated equilibria. The same reasoning naturally applies to the case in which $\delta_E \in (\underline{\delta}_E, \overline{\delta}_E)$.

To illustrate this commitment device, let us consider the doctor-patient example once again. In this case, whenever career concerns are such that TT, MU and TMU may exist, if both MU and TT deliver worst sorting with respect to TMU, and as long as the patient is sufficiently concerned about the future relative to the present, it is optimal for her to pledge to continue to rely on the doctor's services in the future, only if mild treatment is recommended in the current period. This will induce doctors to behave as implied by TMU. On the contrary, consider the case in which getting the right treatment in the present period is of crucial importance. Then, as long as not only TT but also any MU is a better alternative to TMU in terms of discipline, the patient will be better off committing to continue to consult the physician if and only if her condition improves after undergoing the suggested treatment, regardless of whether the doctor recommended more rather than less intensive therapies.

7 Discussion: The Role of Reputation for Ability

As a final result, it is worth noting that informative equilibria would not exist if reputational concerns were only related to preferences. It is the presence of a second dimension of reputation (i.e., reputation for ability) that creates the right incentives for information revelation. To see this, assume that $\alpha = 1$, which implies that there is no uncertainty on ability, and consider a putative informative equilibrium in which the unbiased expert is at least partially revealing her information. This cannot be an equilibrium, since U has a strict incentive to deviate by always sending the message that the biased expert sends less frequently to signal that she is unbiased. This is so precisely because there is no reputational reward of providing a correct evaluation. On the contrary, if we consider putative equilibria in which the unbiased expert always sends a given message regardless of the signal received, these can be informative only if the biased expert partially reveals her information. However, this can never be the case because reputation for ability does not play a role, and B always has a strict incentive to mimic U's strategy. Thus, babbling is the only equilibrium if there is no uncertainty about ability.

This result provides further insight on Morris's (2001) result that reputation can be self-defeating, implying that for high enough reputational concerns of the unbiased expert, information revelation breaks down. Notice, in fact, that our setup is equivalent to assuming that U's reputational concerns are maximum, for the unbiased expert is not concerned at all about current decisions. When we set $\alpha = 1$, as prescribed by Morris, reputational concerns are, in fact, self-defeating. However, our model illustrates that allowing for uncertainty about ability restores the positive value of reputation. Indeed, we find that as long as reputational concerns for ability are present, informative equilibria always exist (for sufficiently high reputational concerns of the biased expert) even when the reputational concerns of the unbiased advisor are greatest.

8 Conclusion

Decision makers often seek the advice of experts before making a decision. The presumption is that an expert has access to valuable information (not available to the decision maker) that is relevant for making correct decisions and that the expert will truthfully report such information to the decision maker. In fact, experts may differ in their abilities to retrieve accurate information and may well have objectives that are not necessarily aligned with those of decision makers.

In the present paper, we analyzed a model of cheap talk where the credibility of the expert's advice hinges upon the decision maker's beliefs about how unbiased and competent the expert is. When the expert and the decision maker interact repeatedly, the expert can use present interaction to affect the beliefs of the decision maker and establish a reputation for being unbiased and competent, thereby increasing the credibility of her future advice.

We show that these reputational concerns on the part of the expert may suffice to achieve truthtelling. However, truthtelling may not necessarily be the outcome preferred by the decision maker. In particular, we highlight the existence of a trade-off between how much the decision maker learns about the expert's ability versus her integrity (i.e., the bias). In particular, with respect to truthtelling, misreporting equilibria lead to more learning about integrity and less about ability. In a dynamic setting in which a decision maker has to make current and future decisions, this trade-off plays an important role. The decision maker may in fact prefer to give up some information on the current state of the world and learn less about the advisor's skills, if learning more about her preferences allows the decision maker to make better decisions in the future.

References

[1] Austen-Smith, D., Fryer, R.G. 2005. An Economic Analysis of “Acting White". *Quarterly Journal of Economics*, 120,551–583.

[2] Bagwell, K. 2007. Signaling and entry deterrence: A multidimensional analysis. *The RAND Journal of Economics*, 38, 670–697.

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

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

[5] Ely, J., Välimäki J. 2003 . Bad reputation. *Quarterly Journal of Economics*, 118(3), 785–813..

[6] Esteban, J., Ray, D. 2006. Inequality, Lobbying, and Resource Allocation. *American Economic Review*, 96, 257–279.

[7] Frankel, A., Kartik, N. 2016. Muddled Information. *working paper*.

[8] Holmstrom, B. 1999. Managerial Incentive Problems: A Dynamic Perspective. *Review of Economic Studies*, 66, 69–182.

[9] Li, W. 2007. Changing One’s Mind when the Facts Change: Incentives of Experts and the Design of Reporting Protocols. *Review of Economic Studies*, 74, 1175–1194.

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

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

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

[13] Scharfstein, D., Stein, J. 1990. Herd Behavior and Investment. *American Economic Review*, 80(3), 465–479.

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

[15] Trueman, B. 1994. Analysis Forecasts and Herding Behavior. *The Review of Financial Studies*, 7(1), 97-194.

A Appendix

A.1 Notation and Terminology

a) $i = U, B$ denotes the preference type of an expert, i.e., unbiased (U) and biased (B).

b) λ_{i,s_1} denotes the probability with which type i reports first-period signal s_1 truthfully. That is, $\lambda_{i,s_1} = \Pr(m_1 = s_1 \mid s_1, i)$.

c) We say that expert i misreports signal s_1 if and only if $\lambda_{i,s_1} < 1$.

d) We say that expert i truthfully reports signal s_1 if and only if $\lambda_{i,s_1} = 1$.

e) The expression misreporting equilibrium denotes an equilibrium in which there exists an $i = U, B$ and a signal $s_1 = 0, 1$ such that $\lambda_{i,s_1} < 1$.

A.2 Characterization of Informative Equilibria

In this section, we characterize the informative equilibria of the game described in Section 2. The game can be solved by backward induction. Without loss of generality, we restrict attention to informative equilibria in which DM interprets message 1 to be (weakly) correlated with signal 1 and hence state 1. We begin by establishing a lemma that will make it easier to analyze the whole game.

Lemma 2 *In any equilibrium in which m_t reveals some information about x_t, DM chooses $a_t(m_t) = m_t$.*

Proof. If m_t is informative about x_t, then $\Pr(x_t = 1 \mid m_t = 0) < \Pr(x_t = 1) < \Pr(x_t = 1 \mid m_t = 1)$. Since $R_t(1,1) = -R_t(1,0)$ and $\Pr(x_t = 1) = \frac{1}{2}$, then $E\left[R_t(a_t = 1, x_t) \mid m_t = 1\right] > E\left[R_t(a_t = 0, x_t) \mid m_t = 1\right]$ and $E\left[R_t(a_t = 0, x_t) \mid m_t = 0\right] > E\left[R_t(a_t = 1, x_t) \mid m_t = 0\right]$. ■

We now proceed by backward induction.

A.2.1 Second Period

Lemma 3 and Lemma 4 below characterize the most informative equilibrium of the second period of the game.

Lemma 3 *In the most informative second period continuation equilibrium: i) B sends $m_2 = 1$ irrespective of s_2; ii) U reports truthfully.*

Proof. In the last period, the expert will not be concerned about her reputation. Thus the biased expert will always claim to have observed signal 1 in order to induce DM to choose action 1. For an unbiased expert with no explicit preferences in favor of a particular action, any strategy is a continuation equilibrium. Without loss of generality we focus on most informative continuation equilibrium in which the unbiased expert acts in the interest of the DM and truthfully reveals her signal. ■

Lemma 4 *At the beginning of the second period, DM retains the incumbent if and only if $V(m_1, x_1) \geq V$ and hires a new expert otherwise.*

Proof. Given lemma 3, it is straightforward to show that:

$$V = \frac{r}{2}\gamma\left(2q - 1\right)$$

$$V(m_1, x_1) = \frac{r}{2}\widehat{\gamma}(m_1, x_1)\left[2\widehat{q}(m_1, x_1) - 1\right]$$

Since both q and $\widehat{q}(m_1, x_1)$ are greater than $\frac{1}{2}$ (i.e. in expectation the expert always has better information than DM), both $V(m_1, x_1)$ and V are strictly positive. Thus, DM always finds it optimal to consult an expert in period 2. In particular, DM will retain the incumbent whenever $V(m_1, x_1) \geq V$ and fire her otherwise. ■

A.2.2 First Period

Assuming that experts and decision makers behave as described by Lemmas 2-4, the continuation payoff of a biased expert at the end of the first period (i.e., when realization (m_1, x_1) has been observed) can be written as $[V(m_1, x_1) + 1]\,\imath(m_1, x_1)$, where

$$\imath(m_1, x_1) = \begin{cases} 1 \text{ if } V(m_1, x_1) \geq V, \\ 0 \text{ otherwise.} \end{cases}$$

Similarly, the continuation payoff of an unbiased expert can be written as $V(m_1, x_1)\imath(m_1, x_1)$.

Hence, for a biased expert who observes signal s_1, the expected continuation payoff of choosing message m_1 reads:

$$\pi_{2,B}(m_1, s_1) = \sum\nolimits_{x_1} \Pr(x_1 \mid s_1) \left[V(m_1, x_1) + 1\right] \imath(m_1, x_1),$$

Similarly, for an unbiased expert who observes s_1, the expected continuation payoff of choosing message m_1 reads:

$$\pi_{2,U}(m_1, s_1) = \sum\nolimits_{x_1} \Pr(x_1 \mid s_1) \left[V(m_1, x_1)\right] \imath(m_1, x_1),$$

Having determined the continuation payoffs, we can write the conditions under which each type of expert has a weak incentive to truthfully reveal a given signal s_1 in the first period. For a biased expert, these conditions read:

$$\delta_E \pi_{2,B}(0,0) - (1 - \delta_E) - \delta_E \pi_{2,B}(1,0) \geq 0 \text{ if } s_1 = 0, \quad (15)$$

$$(1 - \delta_E) + \delta_E \pi_{2,B}(1,1) - \delta_E \pi_{2,B}(0,1) \geq 0 \text{ if } s_1 = 1, \quad (16)$$

For an unbiased expert instead, we have:

$$\pi_{2,U}(0,0) - \pi_{2,U}(1,0) \geq 0 \text{ if } s_1 = 0, \quad (17)$$

$$\pi_{2,U}(1,1) - \pi_{2,U}(0,1) \geq 0 \text{ if } s_1 = 1, \quad (18)$$

We now establish the following lemma that states the properties that an informative equilibrium *cannot* have.

Lemma 5 *An informative equilibrium never satisfies any of the following properties:*
i) m_t *does not reveal information on the state of the world* x_t.
ii) U *always sends* $m_1 = 1$ *regardless of the signal received.*
iii) For some $i = U, B$, $\lambda_{i,s_1} \in (0,1)$ *for every* $s_1 = 0, 1$.
iv) $\lambda_{B,1} \in [0,1)$ *and* $\lambda_{U,1} \in (0,1]$.

Proof. i) We need to show that in an informative equilibrium m_t necessarily reveals some information about x_t for any $t = 1, 2$. In the second period, the only decision relevant information is the information about x_2. Hence, if m_2 did not reveal any information about x_2, the equilibrium could not be informative. In the first period, if m_1 did not reveal information about x_1, the equilibrium could still be informative so long as DM could learn something about either the ability or the integrity of the expert. Note that if m_1 is uninformative about x_1, it is because the expert is sending a message that is independent from the signal received. This implies that DM cannot not learn anything about ability. What about integrity? There are two cases to consider: a) Both U and B follow the same reporting strategy; b) U and B follow different reporting strategies. In the first case, DM obviously learns nothing about integrity, messages are meaningless and the only equilibrium that satisfies these properties is babbling. Hence in this case no decision-relevant learning takes place in period 1. In the second case, there must be a message that is sent more often by U and another message that is sent more often by B. Hence, DM would learn about integrity by observing messages, and would retain more often the expert that reports the message that is sent more often by U. However, this cannot be an equilibrium since B would always deviate to report the message that is sent more often by B.

ii) If this were true, by 5i a necessary condition for the equilibrium to be informative is that B truthfully reports $m_1 = 0$ with positive probability. However, this cannot be part of the equilibrium since reporting $m_1 = 0$ would immediately allow DM to identify the expert as biased and to fire her, providing B with an incentive to always report $m_1 = 1$ (which is the message that also provides current benefits to B).

iii) We first show that it cannot be that both (17) and (18) are satisfied with equality, implying that it cannot be that both $\lambda_{U,1} \in (0,1)$ and $\lambda_{U,0} \in (0,1)$. Note that both (17) and (18) are satisfied with equality if and only if $V(1,1)\iota(1,1) = V(0,1)\iota(0,1)$ and $V(0,0)\iota(0,0) = V(1,0)\iota(1,0)$. Furthermore, note that if both $\lambda_{U,1} \in (0,1)$ and $\lambda_{U,0} \in (0,1)$, then $V(m_1, x_1) > 0$ for all (m_1, x_1).

A trivial case in which both (17) and (18) are satisfied with equality is when $\iota(m_1, x_1) = 0$ for all (m_1, x_1). However this cannot happen if the equilibrium is informative. Indeed, if the equilibrium is informative, there exist some realizations (m_1, x_1) for which $V(m_1, x_1) > V$ and for which it would then be optimal for DM to retain the expert.

We now prove that in all the other cases, if (18) is satisfied with equality then (17) never is. To do so, first consider the case in which $V(1,1)\iota(1,1) = V(0,1)\iota(0,1) > 0$. This case occurs only if $V(1,1) = V(0,1)$. Now note that if the equilibrium is informative,

$\alpha(1,1) > \alpha > \alpha(0,1)$. Hence, $V(1,1) = V(0,1)$ requires that $m_1 = 0$ is a positive signal of integrity so that $\gamma(0,1) > \gamma > \gamma(1,1)$. But if $m_1 = 0$ is a positive signal of integrity, since $\alpha(0,0) > \alpha > \alpha(1,0)$, it is also true that $V(0,0) > V(1,0)$. This implies that $\iota(0,0) = 1$ and $\iota(1,0) = 0$, and hence $V(0,0)\iota(0,0) > V(1,0)\iota(1,0) = 0$. Thus. (17) is always strictly positive.

The only other possible case is when $V(1,1)\iota(1,1) = V(0,1)\iota(0,1) = 0$. This case occurs only if $\iota(1,1) = \iota(0,1) = 0$. Now, if $\iota(1,1) = 0$, it must be that $V(1,1) < V$. Since in an informative equilibrium $\alpha(1,1) > \alpha$, we must have that $m_1 = 1$ is a negative signal of integrity so that $\gamma(1,1) < \gamma$. This implies that $\iota(1,0) = 0$. Hence, if $\iota(0,0) = 0$ we are in the first case analyzed above in which $\iota(m_1, x_1) = 0$ for all (m_1, x_1). Instead, if $\iota(0,0) = 1$, we have $V(0,0)\iota(0,0) > V(1,0)\iota(1,0) = 0$, and again (17) is always strictly positive.

The same line of reasoning applies to show that it cannot be that both (15) and (16) are satisfied with equality implying that it cannot be that both $\lambda_{B,1} \in (0,1)$ and $\lambda_{B,0} \in (0,1)$.

iv) We first show that if $\lambda_{U,1} \in (0,1]$, then it must be that $\lambda_{B,1} = 1$. Given the definition of $\pi_{2,i}(m_1, s_1)$, we have that:

$$\pi_{2,B}(m_1, s_1 = 1) = \pi_{2,U}(m_1, s_1 = 1) + \sum\nolimits_{x_1} \Pr(x_1 \mid s_1 = 1)\imath(m_1, x_1)$$

This implies that the LHS of (16) reads as follows:

$$\begin{aligned} &(1-\delta_E) + \delta_E\left[\pi_{2,B}(1,1) - \pi_{2,B}(0,1)\right] = \\ &= (1-\delta_E) + \delta_E\left[\pi_{2,U}(1,1) - \pi_{2,U}(0,1)\right] + \delta_E \underbrace{\{q\left[\imath(1,1) - \imath(0,1)\right] + (1-q)\left[i(1,0) - i(0,0)\right]\}}_{C} \end{aligned}$$

where $q = \Pr(x_1 = 1 \mid s_1 = 1) > \frac{1}{2}$. If the expression above is strictly positive, then $\lambda_{B,1} = 1$. We now show that $C > 0$ is satisfied whenever (18) is satisfied with equality, which further implies that the expression above is always strictly positive whenever (18) is satisfied with equality. Since $q > \frac{1}{2}$, there are only two cases in which C could be negative:

a) $\iota(1,1) = 0$ and $\iota(0,1) = 1$. Notice that if $\iota(0,1) = 1$, then it must be that $m = 0$ is a positive signal for integrity and hence $m = 1$ a negative one. But then it must be $V(1,0) < V$ and $V(0,0) > V$, and hence $\iota(1,0) = 0$ and $\iota(0,0) = 1$. This implies that $\pi_U(0,1) > \pi_U(1,1)$ which contradicts (18).

b) $\iota(1,1) = \iota(0,1) = \iota(1,0) = 0$ and $\iota(0,0) = 1$. In this case, it is straightforward to notice that $\pi_U(0,1) > \pi_U(1,1)$ which again contradicts (18).

This implies that (16) is satisfied with strict inequality which is equivalent to say that if $\lambda_{U,1} \in (0,1]$, then $\lambda_{B,1} = 1$. ■

A.3 Proof of Proposition 1

We first show that $\left|\widehat{\gamma}^{ME}(m_1,x_1) - \gamma\right| > \left|\widehat{\gamma}^{TT}(m_1,x_1) - \gamma\right| = 0$ for every $m_1 = 0,1$ and $x_1 = 0,1$. Given a realization (m_1,x_1), the update on the prior γ reads:

$$\widehat{\gamma}(m_1,x_1) = \Pr(U \mid m_1,x_1) = \frac{\gamma \Pr(m_1 \mid U,x_1)}{\gamma \Pr(m_1 \mid U,x_1)\Pr(U) + (1-\gamma)\Pr(m_1 \mid B,x_1)}. \quad (19)$$

In a TT equilibrium, both U and B truthfully use the same strategy of truthfully reporting the signal they receive. Since the probability of receiving a given signal is not correlated with the expert's type $i = U,B$, it follows that for any $m_1 = 0,1$ and $x_1 = 0,1$, $\Pr(m_1 \mid U,x_1) = \Pr(m_1 \mid B,x_1)$. Hence $\widehat{\gamma}(m_1,x_1) = \gamma$. This proves that $\left|\widehat{\gamma}^{TT}(m_1,x_1) - \gamma\right| = 0$ for every $m_1 = 0,1$ and $x_1 = 0,1$.

With regard to ME equilibria, we know by Lemma 5(iii) that each type $i = U,B$ can misreport at most one signal. So, let $s' = 0,1$ denote a signal received by the expert. Thanks to Lemma 5(iii) we only need to consider the following three cases:

1) Both U and B report s' truthfully and misreport $1-s'$. First, we show that U and B must misreport $1-s'$ with different probabilities, otherwise the equilibrium is not informative. To see this, suppose the equilibrium is informative and both U and B use the same (misreporting) strategy. If the equilibrium is informative, messages must be correlated with signals. Since signals are in turn correlated with the state of the world, we must have that $\widehat{\alpha}(m_1 = x_1) > \widehat{\alpha}(m_1 \neq x_1)$. At the same time, since both U and B use the same strategy, $\Pr(m_1 \mid U,x_1) = \Pr(m_1 \mid B,x_1)$, and thus $\widehat{\gamma}(m_1,x_1) = \gamma$. All this implies that $\pi_{2,U}(0,0) > \pi_{2,U}(1,0)$ and $\pi_{2,U}(1,1) > \pi_{2,U}(0,1)$. But then (17) and (18) would always be satisfied with strict inequality, implying that U would always truthfully reveal all her signals (contradicting our initial assumption that U misreports). Now, without loss of generality, assume that U reports $1-s'$ with higher probability than B. That is, $\lambda_{U,s'} = \lambda_{B,s'} = 1$ and $\lambda_{B,1-s'} < \lambda_{U,1-s'} < 1$. Since the probability of receiving a given signal is not correlated with the expert's type $i = U,B$, it follows that U reports message

$m_1 = 1-s'$ more frequently than B, and message $m_1 = s'$ less frequently than B. But then, $\Pr(m_1 = 1-s' \mid U, x_1) > \Pr(m_1 = 1-s' \mid B, x_1)$ and $\Pr(m_1 = s' \mid U, x_1) < \Pr(m_1 = s' \mid B, x_1)$. Hence $\widehat{\gamma}(m_1 = 1-s', x_1) > \gamma > \widehat{\gamma}(m_1 = s', x_1)$. Therefore, $\left|\widehat{\gamma}^{ME}(m_1, x_1) - \gamma\right| > 0$ for every $m_1 = 0, 1$ and $x_1 = 0, 1$.

2) U truthfully reports signal s' and misreports signal $1-s'$ while B does the opposite. Lemma 5(iv) implies that if U truthfully reveals signal 1, B must do the same. Hence, in the case under consideration, it must be that $s' = 0$. This implies that $\lambda_{U,0} = 1$, $\lambda_{U,1} < 1$ and $\lambda_{B,0} < 1$, $\lambda_{B,1} = 1$. It is then straightforward to show that $\Pr(m_1 = 0 \mid U, x_1) > \Pr(m_1 = 0 \mid B, x_1)$ and $\Pr(m_1 = 1 \mid U, x_1) < \Pr(m_1 = 1 \mid B, x_1)$ for every $x_1 = 0, 1$. Hence, also in this case we have $\widehat{\gamma}(m_1 = s', x_1) > \gamma > \widehat{\gamma}(m_1 = 1-s', x_1)$. Therefore, also in this case, it is true that $\left|\widehat{\gamma}^{ME}(m_1, x_1) - \gamma\right| > 0$ for every $m_1 = 0, 1$ and $x_1 = 0, 1$.

3) Only one type of expert $i = U, B$ misreports. Without loss of generality, assume that U truthfully reports both s' and $1-s'$, while B truthfully reports s' but misreports $1-s'$ with positive probability. That is, $\lambda_{U,s'} = \lambda_{U,1-s'} = 1$ and $\lambda_{B,s'} = 1$, $\lambda_{B,1-s'} < 1$. But then it is straightforward to show that $\Pr(m_1 = 1-s' \mid U, x_1) > \Pr(m_1 = 1-s' \mid B, x_1)$ and $\Pr(m_1 = s' \mid U, x_1) > \Pr(m_1 = s' \mid B, x_1)$. Hence, $\widehat{\gamma}(m_1 = 1-s', x_1) > \gamma > \widehat{\gamma}(m_1 = s', x_1)$. Therefore, $\left|\widehat{\gamma}^{ME}(m_1, x_1) - \gamma\right| > 0$ for every $m_1 = 0, 1$ and $x_1 = 0, 1$.

We now show that $\left|\widehat{\alpha}^{TT}(m_1, x_1) - \alpha\right| > \left|\widehat{\gamma}^{ME}(m_1, x_1) - \alpha\right|$ for every (m_1, x_1). Let y_{s_1} denote the probability that $m_1 = s_1$ given that the expert has received signal s_1. Note that $y_{s_1} = \gamma\lambda_{U,s_1} + (1-\gamma)\lambda_{B,s_1}$. Now consider the updates on ability when the expert reports a correct message:

$$
\begin{aligned}
\widehat{\alpha}(0,0) &= \frac{\alpha\left[py_0 + (1-p)(1-y_1)\right]}{y_0 q + (1-q)(1-y_1)}, \\
\widehat{\alpha}(1,1) &= \frac{\alpha\left[py_1 + (1-p)(1-y_0)\right]}{y_1 q + (1-q)(1-y_0)},
\end{aligned}
$$

In a truthtelling equilibrium, $y_0 = y_1 = 1$ and hence $\widehat{\alpha}(0,0) = \widehat{\alpha}(1,1) = \frac{\alpha p}{q}$. In a misreporting equilibrium, $y_0 \leq 1$ and $y_1 \leq 1$ with at least one strict inequality. Hence, it is easy to verify that in a misreporting equilibrium $\widehat{\alpha}(0,0) \leq \frac{\alpha p}{q}$ and $\widehat{\alpha}(1,1) \leq \frac{\alpha p}{q}$ with at least one strict inequality. A similar logic applies to show that the same conclusion holds for the cases $\widehat{\alpha}(1,0)$ and $\widehat{\alpha}(0,1)$.

A.4 Proof of Proposition 2

By Lemma 5(iv), there can only be two putative equilibria in which U truthfully reports all her signals:

i) Equilibria in which also B truthfully reports all her signals (*truthtelling equilibria* or TT in short);

ii) Equilibria in which B truthfully reports $s_1 = 1$, and reports $s_1 = 0$ with probability $\lambda_{B,0} < 1$ (*misreporting biased equilibria or* MB in short)

A.4.1 Truthtelling Equilibria (TT)

DM**'s strategy**. Let $\widehat{\gamma}^{TT}(m_1, x_1)$ and $\widehat{\alpha}^{TT}(m_1, x_1)$ denote the value of reputations in a (putative) truthtelling equilibrium. It is straightforward to verify that:

$$\widehat{\gamma}^{TT}(m_1, x_1) = \gamma \text{ for any } (m_1, x_1),$$
$$\underline{\alpha} \equiv \widehat{\alpha}^{TT}(0,1) = \widehat{\alpha}^{TT}(1,0) < \alpha < \widehat{\alpha}^{TT}(1,1) = \widehat{\alpha}^{TT}(0,0) \equiv \overline{\alpha}.$$

The previous updates of α imply that:

$$\underline{q} \equiv \widehat{q}^{TT}(0,1) = \widehat{q}^{TT}(1,0) < q < \widehat{q}^{TT}(1,1) = \widehat{q}^{TT}(0,0) \equiv \overline{q}.$$

Now let $V^{TT}(m_1, x_1)$ denote the value of $V(m_1, x_1)$ in a truthtelling equilibrium. Given the above values of reputations, it is straightforward to show that:

$$\underline{V} \equiv V^{TT}(0,1) = V^{TT}(1,0) < V < V^{TT}(1,1) = V^{TT}(0,0) \equiv \overline{V}. \tag{20}$$

From (20), it follows that in a truthtelling equilibrium DM will retain the incumbent whenever $m_1 = x_1$ and fire her otherwise. Given this retaining strategy, we have that:

$$\imath(m_1, x_1) = \begin{cases} 1 \text{ if } m_1 = x_1, \\ 0 \text{ if } m_1 \neq x_1. \end{cases} \tag{21}$$

B's strategy. By Lemma 5 (iv), we know that if U truthfully reports $s_1 = 1$, then B must truthfully report $s_1 = 1$ too. So, we only need to consider the case in which a biased expert receives $s_1 = 0$. Expression (15) gives the condition for B truthfully report $m_1 = 0$ after observing $s_1 = 0$. By using the expression of B's continuation values, we can write (15) as follows:

$$
\begin{aligned}
&(1-\delta_E)\,a(0) + \delta_E \sum\nolimits_{x_1} \Pr(x_1 \mid s_1 = 0)\,[V(0,x_1)+1]\,\imath(0,x_1) + \\
&-(1-\delta_E)\,a(1) + \delta_E \sum\nolimits_{x_1} \Pr(x_1 \mid s_1 = 0)\,[V(1,x_1)+1]\,i(1,x_1) \geq 0.
\end{aligned} \tag{22}
$$

Now, by using (20), (21) and the fact that $\Pr(x_1 = 0 \mid s_1 = 0) = q$, condition (22) boils down to:

$$
\delta_E \geq \frac{1}{(2q-1)\overline{V} + 2q} \equiv \underline{\delta}_E^{TT}. \tag{23}
$$

U's strategy. We consider the case in which an unbiased expert receives $s_1 = 0$ (a symmetric argument holds for the case in which $s_1 = 1$). Expression (17) gives the condition for U truthfully report $m_1 = 0$ after observing $s_1 = 0$. By using U's continuation values, (17) can be written as:

$$
\sum\nolimits_{x_1} \Pr(x_1 \mid s_1 = 0) V^{TT}(0,x_1) i(0,x_1) - \sum\nolimits_{x_1} \Pr(x_1 \mid s_1 = 0) V^{TT}(1,x_1) i(1,x_1) \geq 0. \tag{24}
$$

Finally, by using (20), (21) and the fact that $\Pr(x_1 = 0 \mid s_1 = 0) = q$, condition (24) simplifies to:

$$
(2q-1)\overline{V} \geq 0, \tag{25}
$$

which is always verified because $q > 1/2$.

Existence intervals with respect to δ_E: A truthtelling equilibrium exists if and only if $\delta_E \in \left[\underline{\delta}_E^{TT}, 1\right]$.

A.4.2 Misreporting Biased Equilibria (MB)

Let $\widehat{\gamma}^{PP}(m_1, x_1)$ and $\widehat{\alpha}^{PP}(m_1, x_1)$ denote the reputation values in a (putative) MB equilibrium. It is straightforward to verify that:

$$\widehat{\gamma}^{MB}(1,1) < \widehat{\gamma}^{MB}(1,0) < \gamma < \widehat{\gamma}^{MB}(0,1) = \widehat{\gamma}^{MB}(0,0),$$
$$\widehat{\alpha}^{MB}(0,1) < \widehat{\alpha}^{MB}(1,0) < \alpha < \widehat{\alpha}^{MB}(1,1) < \widehat{\alpha}^{MB}(0,0).$$

The previous updates of α imply that:

$$\widehat{q}^{MB}(0,1) < \widehat{q}^{MB}(1,0) < q < \widehat{q}^{MB}(1,1) < \widehat{q}^{MB}(0,0).$$

Now let $V^{MB}(m_1, x_1)$ denote the value of $V(m_1, x_1)$ in an MB equilibrium. Given the above values of reputations, it immediately follows that:

$$V^{MB}(1,0) < V < V^{MB}(0,0). \tag{26}$$

In order to prove existence we proceed in two steps:

Step 1) We begin by showing that given U's and B's strategies, DM retains the expert if and only if realizations $(0,0)$ and $(1,1)$ are observed, and fires her after realizations $(0,1)$ and $(1,0)$. In particular, we show that this occurs if and only if $\lambda_{B,0}$ is sufficiently high.

Observe that by condition (26) it immediately follows that DM retains the expert after realization $(0,0)$, and fires her after realization $(1,0)$. Note that this implies that it is indeed a necessary condition for the existence of our putative MB equilibrium that the expert is retained after $(1,1)$. If not, the expert would always be fired when sending $m_1 = 1$. And as a consequence, U (whose only concern is to be retained) would never send $m_1 = 1$ regardless of the signal received (which contradicts her equilibrium strategy).

We now prove that DM retains the expert after $(1,1)$ if and only if $\lambda_{B,0}$ is sufficiently high, and that if the expert is retained after realization $(1,1)$, she must always be fired after realization $(0,1)$. Since $V^{MB}(1,1)$ approaches $\overline{V}$ as the probability of telling the truth approaches one, there always exists a scalar $\lambda'_B \in (0,1)$ such that for $\lambda_{B,0} > \lambda'_B$ the following condition is satisfied:

$$V^{MB}(1,1) \equiv \widehat{\gamma}^{MB}(1,1)(2\widehat{q}^{MB}(1,1) - 1) > \gamma(2q - 1) \equiv V. \tag{27}$$

Similarly, since $V^{MB}(0,1)$ approaches $\underline{V}$ as the probability of telling the truth approaches one, there always exists a scalar $\lambda_B'' \in (0,1)$ such that for $\lambda_{B,0} > \lambda_B''$ the following condition is satisfied:

$$V^{MB}(0,1) \equiv \widehat{\gamma}^{MB}(0,1)(2\widehat{q}^{MB}(0,1) - 1) < \gamma(2q-1) \equiv V. \tag{28}$$

Therefore, to show that whenever the expert is hired in $(1,1)$ she is always fired in $(0,1)$, it is sufficient to show that $\lambda_B' > \lambda_B''$. We proceed as follows: a) We find λ_B''; b) We show that for $\lambda_{B,0} = \lambda_B''$, (27) is never satisfied; c) Since the LHS of (27) is strictly increasing in $\lambda_{B,0}$, we conclude that λ_B' must be strictly greater than λ_B''.

a) λ_B'' is the value of $\lambda_{B,0}$ that satisfies (28) with equality. Substituting the expressions of $\widehat{\gamma}^{MB}(0,1)$ and $\widehat{q}^{MB}(0,1)$ into (28), and solving for the value of $\lambda_{B,0}$ that satisfies this expression with equality, we obtain:

$$\lambda_{B,0}'' = \frac{1}{(1-\gamma)}\left(\frac{(1-p)}{(1-q)} - \gamma\right).$$

b) Using the expressions of $\widehat{\gamma}^{MB}(1,1)$ and $\widehat{q}^{MB}(1,1)$, (27)) can be simplified as follows:

$$\frac{q[p + (1-p)(1-\lambda_{B,0})(1-\gamma)]}{[q + (1-q)(1-\lambda_{B,0})(1-\gamma)]^2} > 1.$$

Substituting $\lambda_{B,0}$ with the expression of λ_B'' that we obtained in part (a), and simplifying we obtain:

$$(p - p^2 - 1/4)(1 - \alpha(1-\alpha)) > 0.$$

Since $(p - p^2 - 1/4) < 0$ for $p > 1/2$ and $(1 - \alpha(1-\alpha)) > 0$, this implies that $(p - p^2 - 1/4)(1 - \alpha(1-\alpha)) < 0$. It follows that (27) is never satisfied for $\lambda_{B,0} = \lambda_B''$.

c) Since the LHS of (27) is strictly increasing in $\lambda_{B,0}$, λ_B' must be strictly greater than λ_B''.

Step 2) We now show that for sufficiently high values of $\lambda_{B,0}$, U's and B's strategies are optimal given DM's strategy.

U's strategy. Let's consider the case in which $s_1 = 1$. U truthfully reports signal $s_1 = 1$ if condition (18) is satisfied. Given DM's strategy, condition (18) becomes:

$$\Pr(x_1 = 1 \mid s_1 = 1)V^{MB}(1,1) \geq \Pr(x_1 = 0 \mid s_1 = 1)V^{MB}(0,0). \tag{29}$$

Now, note that: i) $\Pr(x_1 = 1 \mid s_1 = 1) = q > \Pr(x_1 = 0 \mid s_1 = 1) = 1 - q$; ii) When $\lambda_{B,0} = 0$, we have that $\gamma^{MB}(1,1) = 0$ and $\gamma^{MB}(0,0) = 1$, implying that $V^{MB}(1,1) = 0$ and $V^{MB}(0,0) > 0$; iii) When $\lambda_{B,0} = 1$, we have that $\gamma^{MB}(1,1) = \gamma^{MB}(0,0) = \gamma$, implying that $V^{MB}(1,1) = V^{MB}(0,0) = \overline{V}$; iv) $V^{MB}(1,1)$ and $V^{MB}(0,0)$ are respectively increasing and decreasing in $\lambda_{B,0}$. It then follows that there always exists a scalar $\widetilde{\lambda}_B \in [0,1)$ such that for $\lambda_B \in \left[\widetilde{\lambda}_B, 1\right]$, (29) is satisfied. For the case of $s_1 = 0$, the relevant condition for truthtelling is given by expression (17). It is immediate to note that (17) is always satisfied.

B's strategy. By Lemma 5 (iv), we know that if U truthfully reports $s_1 = 1$, then B must truthfully report $s_1 = 1$ too. Note that B reports signal $s_1 = 0$ with probability $\lambda_{B,0} \in (0,1)$ if and only if condition (15) is satisfied with equality. Given DM's firing strategy, this condition boils down to:

$$\Pr(x_1 = 0 \mid s_1 = 0)\delta_E \left[V^{MB}(0,0) + 1\right] - (1-\delta_E) + \delta_E \Pr(x_1 = 1 \mid s_1 = 0) \left[V^{MB}(1,1) + 1\right] = 0.$$

Using the fact that $\Pr(x_1 = 1 \mid s_1 = 0) = 1 - q$ and $\Pr(x_1 = 0 \mid s_1 = 0) = q$, and rearranging terms, we can write the previous condition as:

$$\delta_E = \frac{1}{[qV^{MB}(0,0)] - (1-q)V^{MB}(1,1)] + 2q} \equiv \delta_E^{MB}(\lambda_{B,0}). \tag{30}$$

Note that since $q > \frac{1}{2}$ and $V^{MB}(0,0) > V^{MB}(1,1)$ for any $\lambda_{B,0} \in (0,1)$, we have that $\delta_E^{MB}(\lambda_{B,0}) \in (0,1)$. Furthermore, since $V^{MB}(1,1)$ and $V^{MB}(0,0)$ are respectively strictly increasing and strictly decreasing in $\lambda_{B,0}$, $\delta_E^{MB}(\lambda_{B,0})$ is strictly increasing in $\lambda_{B,0}$. This allows us to easily identify a lower bound $\underline{\delta}_E^{MB} \in (0,1)$ and an upper bound $\overline{\delta}_E^{MB} \in (0,1)$ such that MB exists if and only if $\underline{\delta}_E^{MB} < \delta_E < \overline{\delta}_E^{MB}$. In particular $\underline{\delta}_E^{MB} \equiv \delta_E^{MB}(\lambda_B^*)$ where $\lambda_B^* = \max\left(\lambda_B', \widetilde{\lambda}_B\right)$, and $\overline{\delta}_E^{MB} \equiv \delta_E^{MB}(1)$. Note further that when $\lambda_{B,0} = 1$, $V^{MB}(0,0) = V^{MB}(1,1) = \overline{V}$ and the RHS of (30) coincides with the RHS of (23). Therefore $\overline{\delta}_E^{MB} = \underline{\delta}_E^{TT}$.

Existence intervals with respect to δ_E: MB can be supported if and only if $\delta_E \in \left[\underline{\delta}_E^{MB}, \underline{\delta}_E^{TT}\right)$ where $\underline{\delta}_E^{MB} \equiv \delta_E^{MB}(\lambda_B^*)$ and $\lambda_B^* = \max\left(\lambda_B', \widetilde{\lambda}_B\right)$.

A.5 Proof of Proposition 3

A necessary and sufficient condition for MB to improve sorting with respect to TT is that $E_0^{MB}(R_2) > E_0^{TT}(R_2)$ or equivalently:

$$[\Pr(0,1 \mid MB) + \Pr(1,0 \mid MB)]\, V + \Pr(0,0 \mid MB) V^{MB}(0,0) + \Pr(1,1 \mid MB) V^{MB}(1,1) >$$
$$[\Pr(0,1 \mid TT) + \Pr(1,0 \mid TT)]\, V + \Pr(0,0 \mid TT) V^{TT}(0,0) + \Pr(1,1 \mid TT) V^{TT}(1,1)$$

Now note that:

- $\Pr(0,1 \mid MB) + \Pr(1,0 \mid MB) = \frac{1}{2} + \frac{1}{2}(2q-1)\left(\gamma\lambda_{B,0} - \lambda_{B,0} - \gamma\right)$
- $\Pr(1,1 \mid MB)\gamma^{MB}(1,1) = \Pr(0,0 \mid MB)\gamma^{MB}(0,0) = \frac{\gamma q}{2}$
- $\Pr(0,1 \mid TT) = \Pr(1,0 \mid TT) = \frac{1-q}{2}$
- $\Pr(0,0 \mid TT) = \Pr(1,1 \mid TT) = \frac{q}{2}$
- $\widehat{q}^{MB}(0,0) = \widehat{q}^{TT}(0,0) = \widehat{q}^{TT}(1,1) = \overline{q}$.

Using these results we can write our condition as follows:

$$\left[\frac{1}{2} + \frac{1}{2}(2q-1)\left(\gamma\lambda_{B,0} - \lambda_{B,0} - \gamma\right)\right] \frac{r}{2}\gamma(2q-1) + \frac{r}{4}q\gamma(2\overline{q}-1) + \frac{r}{4}q\gamma(2\widehat{q}^{MB}(1,1)-1)] >$$
$$(1-q)\frac{r}{2}\gamma\left(2q-1\right) + q\frac{r}{2}\gamma\left(2\overline{q}-1\right)$$

Bringing the first term of the RHS to the LHS, and the second term of the LHS to the RHS, after a bit of algebra, we can rewrite the previous condition as follows:

$$\frac{r}{4}(1-\gamma)(1-\lambda_{B,0})(2q-1)^2 > \frac{r}{2}q\gamma\left[\overline{q} - \widehat{q}^{MB}(1,1)\right]. \tag{31}$$

Notice that the LHS is the expression of the replacement component (10) while the RHS is the expression of the opposite of the continuation component (11). Substituting the equilibrium expressions of $\overline{q}$ and $\widehat{q}^{MB}(1,1)$ into (31), after some algebra, we can rewrite our condition as follows:

$$\alpha^2[(2p-1)(1-x)] + \alpha(2+x) - 1 > 0,$$

where $x \equiv (1-\gamma)(1-\lambda_{B,0})$. Let $\alpha^2[(2p-1)(1-x)] + \alpha(2+x) - 1 \equiv LHS(\alpha, x)$. Note that since $0 < x < 1$, we have that:

a) $LHS(0, x) = -1 < 0$

b) $LHS(1, x) = 2x > 0$

c) $\frac{\partial^2 LHS(\alpha,x)}{\partial \alpha^2} = (2p-1)(1-x) > 0$, i.e. $LHS(\alpha, x)$ is strictly convex in α.

Properties a), b) and c) imply that there exists a unique value $\alpha^*(x)$ such that for $\alpha \in (\alpha^*(x), 1)$, $LHS(\alpha, x) > 0$. In particular, we can find that:

$$\alpha^*(x) = \frac{-(2+x) + [(2+x)^2 + 4(2p-1)(1-x)]^{1/2}}{2(2p-1)(1-x)} \in (0,1)$$

Therefore, we can conclude that for any $\alpha \in (\alpha^*(x), 1)$, $E_0^{MB}(R_2) > E_0^{TT}(R_2)$. Notice that $\alpha^*(x)$ depends on x which in turn depends on $\lambda_{B,0}$, i.e. the probability with which B reports signal zero in an MB equilibrium. Clearly, $\lambda_{B,0}$ must be chosen in the range $\left(\lambda^*_{B,0}, 1\right)$ that is consistent with the existence of MB (as pointed out in the proof of proposition 2).

A.6 Proof of Proposition 4

Lemma 5 (iii) implies that, in an informative equilibrium, U can misreport at most one signal. Hence, we can conveniently divide (putative) equilibria in which U misreports into the following two sub-classes:

i) *Misreporting Unbiased equilibria* (MU): U randomizes after one signal and truthfully reveals the other signal. Since in these equilibria $\lambda_{U,1} \in (0, 1]$, by Lemma 5 (iv), we must have that $\lambda_{B,1} = 1$. All this implies that we can restrict our attention on the existence of the following two putative equilibria belonging to sub-class MU:

- $MU(1)$: U truthfully reports $s_1 = 0$ and randomizes after $s_1 = 1$; B truthfully reports $s_1 = 1$ and reports $s_1 = 0$ with probability $\lambda_{B,0} \in [0, 1]$.

- $MU(0)$: U randomizes after $s_1 = 0$ and truthfully reports $s_1 = 1$; B truthfully reports $s_1 = 1$ and reports $s_1 = 0$ with probability $\lambda_{B,0} \in [0, 1]$.

ii) *Total Misreporting Unbiased equilibria* (TMU): U *lies* about one signal and truthfully reveals the other signal. That is, U always sends the same message independently from the signal observed. We then know by Lemma 5 (ii) that this must be message $m_1 = 0$. Furthermore, we know by Lemma 5 (iii) that in an informative equilibrium B can misreport at most one signal. Hence, we can restrict our attention on the existence of the following two putative equilibria belonging to sub-class TMU:

- $TMU(0)$: U always sends $m_1 = 0$ regardless of the signal received; B truthfully reports $s_1 = 1$ and reports $s_1 = 0$ with probability $\lambda_{B,0} \in [0, 1]$.
- $TMU(1)$: U always sends $m_1 = 0$ regardless of the signal received; B truthfully reports $s_1 = 0$ and reports $s_1 = 1$ with probability $\lambda_{B,1} \in [0, 1]$.

We now prove the existence of each of the equilibria outlined above.

A.6.1 $MU(1)$ Equilibria

We first prove that there exist $MU(1)$ equilibria where $\lambda_{B,0} = 1$ (i.e., $MU(1)$ equilibria where B truthfully reports both signals). We then move on to prove that there also exist $MU(1)$ equilibria where $\lambda_{B,0} < 1$ (i.e., $MU(1)$ equilibria where B truthfully reports $s_1 = 1$ and misreports $s_1 = 0$).

Case in which B truthfully reports both $s_1 = 1$ and $s_1 = 0$. Let $\widehat{\alpha}^{MU(1)}(m_1, x_1)$ and $\widehat{\gamma}^{MU(1)}(m_1, x_1)$ denote the value of reputations in this (putative) $MU(1)$ equilibrium. It is straightforward to verify that:

$$\widehat{\alpha}^{MU(1)}(1,0) < \widehat{\alpha}^{MU(1)}(0,1) < \alpha < \widehat{\alpha}^{MU(1)}(0,0) < \widehat{\alpha}^{MU(1)}(1,1),$$
$$\widehat{\gamma}^{MU(1)}(1,1) = \widehat{\gamma}^{MU(1)}(1,0) < \gamma < \widehat{\gamma}^{MU(1)}(0,0) < \widehat{\gamma}^{MU(1)}(0,1).$$

Given the above values of reputations, we have that:

$$V^{MU(1)}(1,0) < V < V^{MU(1)}(0,0). \tag{32}$$

In order to prove existence we proceed in two steps.

Step 1) We show that given U's and B's strategies, DM retains the expert if and only if realizations $(0,0)$ and $(1,1)$ are observed. In particular, we show that this occurs if and only if $\lambda_{U,1}$ is sufficiently high. First, note that condition (32) implies that DM retains the expert after $(0,0)$ and fires the expert after $(1,0)$. This also implies that a necessary condition for the existence of our equilibrium is that the expert is retained after $(1,1)$. Indeed, if this did not occur, the expert would always be fired after sending $m_1 = 1$, and hence U (whose concern is to be retained) would never send $m_1 = 1$ (which contradicts U's equilibrium strategy).

We now show that DM retains the expert after $(1,1)$ if and only if $\lambda_{U,1}$ is sufficiently high. Note, that DM retains the expert after $(1,1)$ if and only if the following condition is satisfied:

$$\widehat{\gamma}^{MU(1)}(1,1)(2\widehat{q}^{MU(1)}(1,1)-1) > \gamma(2q-1). \tag{33}$$

Substituting the equilibrium values of $\widehat{\gamma}^{MU(1)}(1,1)$ and $\widehat{q}^{MU(1)}(1,1)$ and solving (33) for $\lambda_{U,1}$ we obtain:

$$\lambda_{U,1} > \frac{q-\gamma q}{p-\gamma q} \equiv \lambda_{U,1}^{'} \in (0,1)\,.$$

Hence, condition (33) is satisfied - and DM retains the expert after $(1,1)$ - if and only if $\lambda_{U,1} > \lambda'_{U,1}$.

We now show that when $\lambda_{U,1} > \lambda'_{U,1}$, DM must fire the expert after $(0,1)$. Note that DM fires the expert after $(0,1)$ if and only if the following condition is satisfied:

$$\widehat{\gamma}^{MU(1)}(0,1)(2\widehat{q}^{MU(1)}(0,1)-1) < \gamma(2q-1). \tag{34}$$

Substituting the equilibrium values of $\widehat{\gamma}^{MU(1)}(0,1)$ and $\widehat{q}^{MU(1)}(0,1)$ into (34) and simplifying, (34) becomes:

$$\frac{[(1-q)+q(1-\lambda_U)][(1-p)+p\gamma(1-\lambda_U)]}{[(1-q)+q\gamma(1-\lambda_U)]^2} - 1 < 0.$$

If we now substitute $\lambda_{U,1}$ with the closed form solution $\lambda'_{U,1}$ obtained above, the last inequality boils down to:

$$(2pq-q^2-p^2)(p-\gamma q) < 0.$$

Since $(2pq-q^2-p^2) < 0$ and $(p-\gamma q) > 0$, this last inequality is always satisfied. Hence condition (34) is always satisfied when $\lambda_{U,1} = \lambda'_{U,1}$. Now note that the LHS of (34) is strictly decreasing in $\lambda_{U,1}$ while the RHS does not depend on $\lambda_{U,1}$. Hence we can conclude that condition (34) is satisfied - and thus DM fires the expert after $(0,1)$ - for any $\lambda_{U,1} > \lambda'_{U,1}$.

Step 2) We now show that U's and B's strategies are optimal given DM's strategy outlined in Step 1 and given the constraint $\lambda_{U,1} \geq \lambda'_{U,1}$. First, note that by Lemma 5 (iii), U will always report signal $s_1 = 0$ truthfully if she misreports signal $s_1 = 1$. Second, we know by lemma 5 (iv) that if U reports $s_1 = 1$ with positive probability, B will report $s_1 = 1$ truthfully. Hence, there are only two conditions that we must show that are satisfied in our $MU(1)$ equilibrium. The first one is the condition that makes sure that U randomizes when receiving $s_1 = 1$, that is:

$$qV^{MU(1)}(1,1) = (1-q)V^{MU(1)}(0,0). \tag{35}$$

The second one is the condition that makes sure that B truthfully reports $s_1 = 0$, which can be written as:

$$\delta_E[qV^{MU(1)}(0,0) - (1-q)V^{MU(1)}(1,1) + 2q - 1] > 1 - \delta_E. \tag{36}$$

Note that since $q > \frac{1}{2}$, if condition (35) is satisfied, then it must be that $qV^{MU(1)}(0,0) > (1-q)V^{MU(1)}(1,1)$, which in turn guarantees that the LHS of (36) is strictly increasing in δ_E. Since the RHS is always strictly decreasing in δ_E, we can conclude that if condition (35) is satisfied, then there always exists a value of δ_E above which (36) is satisfied as well. This means we only need to show that condition (35) is indeed satisfied for some $\lambda_{U,1} \in \left(\lambda'_{U,1}, 1\right)$. Note that:

(i) If $\lambda_{U,1} = \lambda'_{U,1}$, $V^{MU(1)}(1,1) = V < V^{MU(1)}(0,0)$. Hence, if α is sufficiently small (so that q is sufficiently small too), the LHS of (35) is smaller than the RHS;

(ii) If $\lambda_{U,1} = 1$, $V^{MU(1)}(1,1) = V^{MU(1)}(0,0)$ and the LHS of (35) is larger than the RHS.

Therefore, by continuity, as long as α is sufficiently small, there always exists an $\lambda_{U,1} \in \left(\lambda'_{U,1}, 1\right)$ such that condition (35) is satisfied.

Case in which B truthfully reports $s_1 = 1$ and misreports $s_1 = 0$. First note that the chain of inequalities given by (32) holds true for any $\lambda_{B,0} < 1$. Hence, DM retains the expert after $(0,0)$ and fires her after $(1,0)$. Furthermore, we know by Lemma 5 parts (iii) and (iv) that if $\lambda_{U,1} \in (0,1)$, then it must be that $\lambda_{U,0} = 1$ and $\lambda_{B,1} = 1$. Hence, we only need to prove that there exist a $\lambda_{B,0} \in (0,1)$ and a $\lambda_{U,1} \in (0,1)$ such that the following three conditions are simultaneously satisfied:

$$qV^{MU(1)}(1,1) = (1-q)V^{MU(1)}(0,0), \tag{37}$$

$$\delta_E q[V^{MU(1)}(0,0)+1] = (1-\delta_E) + \delta_E(1-q)[V^{MU(1)}(1,1)+1], \tag{38}$$

$$V^{MU(1)}(0,1) < V < V^{MU(1)}(1,1). \tag{39}$$

Condition (37) is the condition that must be satisfied for U to randomize after $s_1 = 1$. Condition (38) is the condition that must be satisfied in order for B to randomize after $s_1 = 0$. Finally, condition (39) is the condition that must be satisfied in order for DM to retain the expert after $(1,1)$ and fire the expert after $(0,1)$.

First, let's consider condition (37). Let $\lambda^*_{U,1} \in (0,1)$ be the value of $\lambda_{U,1}$ that satisfies (37) when $\lambda_{B,0} = 1$ (we know by the proof of the case in which B reports truthfully that $\lambda^*_{U,1}$ exists). Now note that $V(1,1)$ is strictly increasing in $\lambda_{B,0}$ while $V(0,0)$ is strictly decreasing in $\lambda_{B,0}$. Hence, when $\lambda_{B,0} = 1-\varepsilon$ (where $\varepsilon > 0$) and $\lambda_{U,1} = \lambda^*_{U,1}$ we have that: $qV^{MU(1)}(1,1) < (1-q)V^{MU(1)}(0,0)$. By the proof of proposition 2 we also know that when $\lambda_{B,0} = 1-\varepsilon$ and $\lambda_{U,1} = 1$ (i.e., when we are in an MB equilibrium) we have that: $qV^{MU(1)}(1,1) > (1-q)V^{MU(1)}(0,0)$. But then, when $\lambda_{B,0} = 1-\varepsilon$, by continuity there must exist a $\lambda_{U,1} \in (\lambda^*_U, 1)$ such that $qV^{MU(1)}(1,1) = (1-q)V^{MU(1)}(0,0)$.

Second, let's consider condition (39). We know by the proof of the case in which B truthfully reports that when $\lambda_{B,0} = 1$ and $\lambda_{U,1} = \lambda^*_{U,1}$, we have that $V^{MU(1)}(0,1) < V < V^{MU(1)}(1,1)$. Note that when $\lambda_{B,0} = 1-\varepsilon$ and $\lambda_{U,1} = \lambda^*_{U,1}$, the previous inequality is still satisfied by continuity (since $V^{MU(1)}(1,1)$ is strictly increasing in $\lambda_{B,0}$, and $V^{MU(1)}(0,1)$ strictly decreasing in $\lambda_{B,0}$, ε must be chosen small enough to ensure that this inequality holds true). Finally, note that when $\lambda_{B,0} = 1-\varepsilon$ and $\lambda_{U,1} \in \left(\lambda^*_{U,1}, 1\right)$, the inequality above holds a fortiori because $V^{MU(1)}(1,1)$ is strictly increasing in $\lambda_{U,1}$ and $V^{MU(1)}(0,1)$ is strictly decreasing in $\lambda_{U,1}$.

Finally, let's consider condition (38). If we solve it for δ_E we obtain:

$$\delta_E = \frac{1}{qV^{MU(1)}(0,0) - (1-q)V^{MU(1)}(1,1) + 2q}. \tag{40}$$

If (37) is satisfied, $qV(0,0) - (1-q)V(1,1)$ is strictly greater than zero and hence the denominator is strictly greater than one, which in turn implies that the RHS is always larger than zero and smaller than one. Hence, we can conclude that given a value of $\lambda_{B,0} \in (0,1)$ and $\lambda_{U,1} \in (0,1)$ for which (37) and (39) are satisfied, we can always find a value of $\delta_E \in (0,1)$ that guarantees that condition (38) is satisfied too.

Existence intervals with respect to δ_E. Given the analysis of the two cases above, by continuity we can conclude that a $MU(1)$ equilibrium exists for $\delta_E \in [\underline{\delta}_E^{MU(1)}, 1]$ where $\underline{\delta}_E^{MU(1)}$ is the smallest value that the RHS of (40)) takes in $MU(1)$.

A.6.2 $MU(0)$ Equilibria

Also for this case, we first prove that there exist $MU(0)$ equilibria where $\lambda_{B,0} = 1$, and then move on to prove that there also exist $MU(0)$ equilibria where $\lambda_{B,0} \in [0,1)$.

Case in which B truthfully reports both $s_1 = 1$ and $s_1 = 0$. Let $\widehat{\alpha}^{MU(0)}(m_1, x_1)$ and $\widehat{\gamma}^{MU(0)}(m_1, x_1)$ denote the value of reputations in this (putative) $MU(0)$ equilibrium. It is straightforward to verify that:

$$\widehat{\alpha}^{MU(0)}(1,0) < \widehat{\alpha}^{MU(0)}(0,1) < \alpha < \widehat{\alpha}^{MU(0)}(0,0) < \widehat{\alpha}^{MU(0)}(1,1),$$
$$\widehat{\gamma}^{MU(0)}(0,0) = \widehat{\gamma}^{MU(0)}(0,1) < \gamma < \widehat{\gamma}^{MU(0)}(1,1) < \widehat{\gamma}^{MU(0)}(1,0).$$

Given the above values of reputations, we have that $V^{MU(0)}(1,1) > V > V^{MU(0)}(0,1)$. Therefore, DM retains the incumbent after observing $(1,1)$ and fires the incumbent after observing $(0,1)$. But then, a necessary condition for the existence of the equilibrium is that DM retains the incumbent after $(0,0)$. If not, the expert would always be fired when sending message zero and hence an unbiased expert would never send $m_1 = 0$ (what contradicts her equilibrium strategy). Therefore, existence requires that:

$$V^{MU(0)}(0,0) > V. \tag{41}$$

By applying the same line of reasoning we used to prove the existence of $MU(1)$ equilibria in which B reports truthfully, we can show that: i) condition (41) is satisfied if and only if $\lambda_{U,0} > \frac{q-\gamma q}{p-\gamma q} \equiv \lambda'_{U,0}$; For $\lambda_{U,0} > \lambda'_{U,0}$, we also have that DM fires the expert after $(1,0)$; iii) U's and B's equilibrium strategies are optimal given DM's retaining strategy and the constraint $\lambda_{U,0} > \lambda'_{U,0}$. Hence, also these $MU(0)$ equilibria are characterized by DM retaining the expert after $(0,0)$ and $(1,1)$, and firing her after $(1,0)$ and $(0,1)$, and by U lying with a sufficiently small probability. We also note that, as in the case of $MU(1)$, δ_E must be above a certain threshold in order for B's behavior to be consistent with the equilibrium. In particular, condition (15) must be satisfied with strict inequality for this to hold. Since the equilibrium behavior of U implies that $qV^{MU(0)}(0,0) = (1-q)V^{MU(0)}(1,1)$, (15) boils down to $\delta_E(2q-1) - 1 - \delta_E > 0$, which in turn implies that $\delta_E > \frac{1}{2q}$.

Case in which B truthfully reports $s_1 = 1$ and misreports $s_1 = 0$. Existence can be proved by applying the same line of reasoning we used to prove the existence of $MU(1)$ equilibria in which B truthfully reports $s_1 = 1$ and randomizes after $s_1 = 0$. Here we note that an $MU(0)$ equilibrium in which both B and U misreport $s_1 = 0$ must be characterized by $\lambda_{U,0} < \lambda_{B,0}$. To see this, consider that for U to misreport $s_1 = 0$, (17) must be satisfied with equality, that is:

$$qV^{MU(0)}(0,0) = (1-q)V^{MU(0)}(1,1).$$

Since $q > 1-q$, the only way to have equality is that $V^{MU(0)}(1,1) > V^{MU(0)}(0,0)$. Now note that in the equilibrium under consideration $\widehat{\alpha}(0,0) > \widehat{\alpha}(1,1)$. Hence, to have that $V^{MU(0)}(1,1) > V^{MU(0)}(0,0)$, it must be that $\widehat{\gamma}(1,1) > \widehat{\gamma}(0,0)$. By proposition 1 this can occur only if U sends message 1 more often than B. Being $\lambda_{U,1} = \lambda_{B,1} = 1$, it must then be that $\lambda_{U,0} < \lambda_{B,0}$.

Finally we note that, based on the analysis above of $MU(0)$ equilibria in which B truthfully reports both signals, we obtain that $MU(0)$ equilibria in which B misreports exist for $\delta_E = \frac{1}{2q}$.

Existence intervals with respect to δ_E Given the analysis of the two cases above, we can conclude that $MU(0)$ equilibria exist for $\delta_E \in [\underline{\delta}_E^{MU(0)}, 1]$, where $\underline{\delta}_E^{MU(0)} = 1/2q$.

A.6.3 TMU Equilibria

We first show the existence of TMU equilibria in which U sends $m_1 = 0$ regardless of s_1, and B truthfully reports both her signals. This will prove the existence of $TMU(0)$ equilibria with $\lambda_{B,0} = 1$, as well as the existence of $TMU(1)$ equilibria with $\lambda_{B,1} = 1$. In what follows, we will denote these equilibria in which B reports truthfully by simply using the upper-script TMU.

First of all, it is straightforward to verify that:

$$\underline{\alpha} \equiv \widehat{\alpha}^{TMU}(1,0) < \widehat{\alpha}^{TMU}(0,1) < \alpha < \widehat{\alpha}^{TMU}(0,0) < \widehat{\alpha}^{TMU}(1,1) = \overline{\alpha}, \tag{42}$$

$$0 = \widehat{\gamma}^{TMU}(1,0) = \widehat{\gamma}^{TMU}(1,1) < \gamma < \widehat{\gamma}^{TMU}(0,0) = \widehat{\gamma}^{TMU}(0,1). \tag{43}$$

Given the above values of reputations, we have that:

$$0 = V^{TMU}(1,0) = V^{TMU}(1,1) < V < V^{TMU}(0,0), \tag{44}$$

$$V \gtrless V^{TMU}(0,1). \tag{45}$$

*DM***'s strategy**. From (44) it follows that DM will always retain the expert whenever $(m_1, x_1) = (0,0)$ and always fire her when $(m_1, x_1) = (1,0), (1,1)$. (45) highlights that DM possibly retains the expert also when $(m_1, x_1) = (0,1)$. In what follows we will focus on the case in which the expert is fired after $(0,1)$.[22] This occurs if $V > V^{TMU}(0,1)$, which holds true so long as α, γ and p satisfy the following inequality:

$$\gamma < \frac{(1-q)^2 - (1-p)}{q^2} \equiv \gamma'(\alpha, p). \tag{46}$$

Since $\gamma \in (0,1)$, we must have that $0 < \gamma'(\alpha,p) < 1$. First, note that $\gamma'(\alpha,p) < 1$ for all values of $\alpha \in (0,1)$ and $p \in \left(\frac{1}{2}, 1\right)$. Second, note that $\gamma'(\alpha,p) > 0$ if and only if $p > \frac{-1+\alpha+\alpha^2+\sqrt{1-2\alpha+2\alpha^2}}{2\alpha^2} \equiv p'$, and that $p' \in \left(\frac{3}{4}, 1\right)$ for all values of $\alpha \in (0,1)$. Summing up,

[22] Following a similar line of reasoning, one can show that there exist also TMU equilibria in which DM retains the expert after $(0,1)$.

DM fires the expert after $(0,1)$ if and only if $p > p'$ and γ satisfies (46).

U's strategy. Given DM's strategy above, the two conditions that must hold for U to send $m_1 = 0$ when she receives $s_1 = 1$ and $s_1 = 0$ read respectively:

$$0 < (1-q)V^{TMU}(0,0),$$

$$qV^{TMU}(0,0) > 0.$$

It is immediate to see that the previous conditions are always satisfied.

B's strategy. Given DM's strategy, in order for B to truthfully report both $s_1 = 1$ and $s_1 = 0$, the two following conditions must be satisfied:

$$(1-\delta_E) \geq \delta_E(1-q)\left[V^{TMU}(0,0)+1\right], \tag{47}$$

$$\delta_E q\left[V^{TMU}(0,0)+1\right] \geq (1-\delta_E). \tag{48}$$

It is easy to show that both conditions are simultaneously satisfied for intermediate values of δ_E, namely for:

$$\frac{1}{qV^{TMU}(0,0)+1+q} \leq \delta_E \leq \frac{1}{(1-q)V^{TMU}(0,0)+2-q}.$$

Note that since $\frac{1}{2} < q < 1$ both the LHS and the RHS take values that are strictly between zero and one.

This completes the proof of the existence of $TMU(0)$ and $TMU(1)$ equilibria in which B truthtells. We now complete the proof by considering the cases in which B misreports one of her signals.

$TMU(0)$ equilibria where B truthfully reports $s_1 = 1$, and reports $s_1 = 0$ with probability $\lambda_{B,0} \in (0,1)$ By continuity, the chains of inequalities given by (42), (43) and hence (44) (which all held true for $\lambda_{B,0} = 1$) continue to hold so long as $\lambda_{B,0}$ is sufficiently close to 1. The same applies to condition $V > V^{TMU(0)}(0,1)$ which (similarly to the case in which $\lambda_{B,0} = 1$) further requires that γ be sufficiently small and p sufficiently large (note that the thresholds values of γ and p are respectively smaller and larger than the threshold values of the case in which $\lambda_{B,0} = 1$). All this implies that as long as $\lambda_{B,0}$ and p are sufficiently large, and γ is sufficiently small, both DM's strategy and U's strategy are the same as in TMU.

In order for B to randomize after $s_1 = 0$ we must have that condition (48) now holds with equality. It is immediate to verify that this occurs when $\delta_E = \frac{1}{qV^{TMU(0)}(0,0)+1+q}$. Finally, note that when (48) is satisfied with equality, (47) is satisfied with strict inequality implying that B truthfully reports $s_1 = 1$.

We conclude by noticing that when $\lambda_{B,0} = 0$ we have an equilibrium in which U sends $m_1 = 0$ regardless of her signal, and B sends $m_1 = 1$ regardless of her signal. Hence, no information is revealed about x_1. We know by lemma 5(i) that this cannot be an informative equilibrium. Hence, it cannot be that $\lambda_{B,0} = 0$.

$TMU(1)$ equilibria where B truthfully reports $s_1 = 0$ and reports $s_1 = 1$ with probability $\lambda_{B,1} \in (0,1)$ By continuity, the chains of inequalities given by (42), (43) and hence (44) (which all held true for $\lambda_{B,1} = 1$) continue to hold so long as $\lambda_{B,1}$ is sufficiently large. The same applies to condition $V > V^{TMU(1)}(0,1)$ which (similarly to the case in which $\lambda_{B,1} = 1$) further requires that γ be sufficiently small and p sufficiently large (note that the thresholds values of γ and p are respectively smaller and larger than the threshold values of the case in which $\lambda_{B,1} = 1$). All this implies that as long as $\lambda_{B,1}$ and p are sufficiently large and γ is sufficiently small, both DM's strategy and U's strategy are the same as in TMU.

In order for B to randomize after $s_1 = 1$, we must have that (47) holds with equality. It is immediate to verify that this occurs when $\delta_E = \frac{1}{(1-q)V^{TMU(1)}(0,0)+2-q}$. Finally, note that when (47) is satisfied with equality, (48) is satisfied with strict inequality implying that B truthfully reports $s_1 = 0$.

We conclude by noticing that when $\lambda_{B,1} = 0$ we have an equilibrium in which both U and B send $m_1 = 0$ regardless of their signals and thus no information is revealed about x_1. We know by lemma 5(i) that this cannot be an informative equilibrium. Hence, it cannot be that $\lambda_{B,1} = 0$.

Existence intervals with respect to δ_E Given the analysis of the cases above, by continuity we can conclude that TMU exists for $\delta_E \in [\underline{\delta}_E^{TMU}, \overline{\delta}_E^{TMU}]$, where $\underline{\delta}_E^{TMU}$ is the smallest value that expression $\frac{1}{qV^{TMU}(0,0)+1+q}$ takes in a TMU equilibrium; and $\overline{\delta}_E^{TMU}$ is the largest value that expression $\frac{1}{(1-q)V^{TMU)}(0,0)+2-q}$ takes in a TMU equilibrium.

A.7 Proof of Proposition 5

We prove proposition 5 for MU equilibria in which U truthfully reports $s_1 = 0$ and randomizes after $s_1 = 1$ (i.e., those equilibria that we denoted with $MU(1)$ in the proof of proposition 4). The same line of reasoning applies to show that the results extend to MU equilibria in which U truthfully reports $s_1 = 1$ and randomizes after $s_1 = 0$ (i.e., those equilibria that we denoted with $MU(0)$ in the proof of proposition 4).

A necessary condition for $MU(1)$ equilibria to improve sorting with respect to TT equilibria is that $E_0^{MU(1)}(R_2) - E_0^{TT}(R_2) > 0$. This condition can be equivalently written as:

$$\begin{aligned} & [\Pr(0,1 \mid MU(1)) + \Pr(1,0MU(1)) - \Pr(0,1 \mid TT) - \Pr(1,0 \mid TT)]\, V + \\ + \quad & \Pr(1,1 \mid MU(1))V^{MU(1)}(1,1) + \Pr(0,0 \mid MU(1))V^{MU(1)}(0,0) + \\ - \quad & \Pr(1,1 \mid TT)V^{TT}(1,1) + \Pr(0,0 \mid TT)V^{TT}(0,0) > 0 \end{aligned}$$

We now show that this previous inequality is never satisfied. Note that:

- $\Pr(1,1 \mid MU(1))\gamma^{MU(1)}(1,1) = \frac{\gamma}{2}q\lambda_{U,1}$.
- $\Pr(0,0 \mid MU(1))\gamma^{MU(1)}(0,0) = \frac{\gamma}{2}(1-(1-q)\lambda_{U,1})$.
- $\widehat{q}^{MU(1)}(1,1) = \widehat{q}^{TT}(1,1) = \overline{q}$.

Thus, we can write the previous inequality as:

$$\begin{aligned} & [\Pr(0,1 \mid MU(1)) + \Pr(1,0MU(1)) - \Pr(0,1 \mid TT) - \Pr(1,0 \mid TT)]\, \gamma(2q-1) + \\ + \quad & \frac{\gamma}{2}(1-(1-q)\lambda_{U,1})(2q^{MU(1)}(0,0)-1) + \frac{\gamma}{2}q\lambda_{U,1}(2\overline{q}-1) - q\gamma(2\overline{q}-1) > 0 \end{aligned}$$

Now note that:

$$[\Pr(0,1 \mid MU(1)) + \Pr(1,0MU(1)) - \Pr(0,1 \mid TT) - \Pr(1,0 \mid TT)]\,\gamma(2q-1) =$$
$$= \quad -\frac{1}{2}(1-2p)^2\alpha^2\gamma(1-\lambda_{U,1})$$

Using this result, and the expressions of $q^{MU(1)}(0,0)$ and $\overline{q}$, after a bit of algebra our condition can be written as:

$$\frac{AB}{C} > 0$$

where:

- $A = (1-2p)^2\,\alpha\gamma\,(1-\lambda_{U,1})$
- $B = -\gamma(2-\lambda_U)+\alpha(-2p+\gamma+2p\gamma+\gamma^2+\gamma\sigma_{G1}-2p\gamma\lambda_{U,1}-\gamma^2\lambda_{U,1})+\alpha^2\left[(\gamma\,(2p-1)\left[1-\gamma\,(1-\lambda_{U,1})\right]\right]$
- $C = 2\left\{1+\gamma(1-\lambda_{U,1})+\alpha\,(2p-1)\left[1-\gamma\,(1-\lambda_{U,1})\right]\right\}$

It is easy to verify that $A > 0$ and $C > 0$. Thus, our condition is satisfied whenever $B > 0$. Note that B is quadratic in α. In particular:

a) When $\alpha = 0$, $B = -\gamma(2-\lambda_U) < 0$;

b) When $\alpha = 1$, $B = -2(1-\gamma)\left[p+\gamma(1-p)-\gamma\lambda_U(1-p)\right] < 0$;

c) $\frac{\partial^2 B}{\partial\alpha^2} = 2\gamma\,(2p-1)\left[1-\gamma\,(1-\lambda_U)\right] > 0$, i.e. is strictly convex.

Note that a), b) and c) imply that $B < 0$ for any $\alpha \in (0,1)$. Hence inequality $\frac{AB}{C} > 0$ is never satisfied.

A.8 Proof of Proposition 6

In order to find an instance in which TMU may improve welfare with respect to TT, we consider the equilibrium in which the expert is hired only after $(0,0)$. Since it is straightforward that discipline is always worst in TMU, a necessary condition for TMU to improve welfare with respect to TT is that it must improve sorting. That is, we need that

$E_0[R_2 \mid TMU] > E_0[R_2 \mid TT]$ or equivalently:

$$\left\{ \begin{array}{l} [\Pr(0,1 \mid TMU) + \Pr(1,0 \mid TMU) + \Pr(1,1 \mid TMU)]\, V + \\ + \Pr(0,0 \mid TMU) V^{TMU}(0,0) \end{array} \right\} >$$
$$\left\{ \begin{array}{l} [\Pr(0,1 \mid TT) + \Pr(1,0 \mid TT)]\, V + \\ + \Pr(1,1 \mid TT) V^{TT}(1,1) + \Pr(0,0 \mid TT) V^{TT}(0,0) \end{array} \right\}.$$

Now, note that:

- $\Pr(0,1 \mid TMU) + \Pr(1,0 \mid TMU) + \Pr(1,1 \mid TMU) = \frac{1}{4}\left[3 - \gamma + (2q-1)(-1+\gamma)\right];$
- $\Pr(0,0 \mid TMU) V^{TMU}(0,0) = V^{TMU}(0,0) = \frac{\gamma}{2}(2q_{00}^{TMU} - 1);$
- $V^{TT}(1,1) = V^{TT}(0,0) = \gamma(2\overline{q} - 1);$
- $\Pr(1,1 \mid TT) = \Pr(0,0 \mid TT) = q.$

Hence, we can write the last inequality reads:

$$\begin{array}{l} \frac{1}{4}\left[3 - \gamma + (2q-1)(-1+\gamma)\right] \gamma\,(2q-1) + \\ + \frac{\gamma}{2}(2q_{00}^{TMU} - 1) > (1-q)\gamma(2q-1) + q\gamma(2\overline{q} - 1). \end{array}$$

Using the equilibrium values of q_{00}^{TMU} and $\overline{q}$, and then simplifying, we obtain the following equivalent condition:

$$q - \gamma(1-q) > 2p - \frac{p + \gamma(1-p)}{q + \gamma(1-q)}. \tag{49}$$

Notice that (49) has the following properties: i) the LHS (RHS) is strictly decreasing (increasing) in γ; ii) When $\gamma = 0$, $LHS = q > RHS = \frac{p(2q-1)}{q}$ for all $\alpha \in (0,1)$ and $p \in (\frac{1}{2}, 1)$;iii) When $\gamma = 1$, $LHS = 2q - 1 < RHS = 2p - 1$ for all $\alpha \in (0,1)$ and $p \in (\frac{1}{2}, 1)$. Hence, for all $\alpha \in (0,1)$ and $p \in (\frac{1}{2}, 1)$, there always exist a threshold $\overline{\gamma}(\alpha, p) \in (0,1)$ such that for $\gamma < \overline{\gamma}(\alpha, p)$ (49) is satisfied.

Now, the equilibrium under consideration is a TMU equilibrium where DM fires the expert after $(0,1)$. We know from the proof of Proposition 5 that the existence of such an equilibrium requires that $p > p'$ and that $\gamma < \gamma'(\alpha, p)$. So, let us define $\gamma^{TMU} =$

$\min[\overline{\gamma}(\alpha, p), \gamma'(\alpha, p)]$. Then the following is true: For all $\alpha \in (0, 1)$, $p > p'$ and $\gamma < \gamma^{TMU}$, there always exists a TMU equilibrium that improves sorting with respect to TT. This completes the proof.

A.9 Proof of Proposition 7

Note that by (17) and (18), δ_E does not affect the behavior of U. Hence we can focus on the behavior of B.

The following points allow us to complete the proof.

1) Based on the proof of proposition 2, we know that MB and TT never coexist, and that MB exists for $\delta_E \in (\underline{\delta}_E^{MB}, \underline{\delta}_E^{TT})$ where $\underline{\delta}_E^{MB} < \underline{\delta}_E^{TT}$.

2) Based on the proof of proposition 4, by (18) if U misreports on $s_1 = 1$ it must be that $V^{MU(1)}(0,0) > \overline{V} > V^{MU(1)}(1,1)$, and by (17) if she misreports after $s_1 = 0$ it must be that $V^{MU(0)}(1,1) > \overline{V} > V^{MU(0)}(0,0)$. Using the expressions for $\underline{\delta}_E^{TT}$, $\underline{\delta}_E^{MU(0)}$ and $\underline{\delta}_E^{MU(1)}$ defined in the proofs of Propositions 2 and 4, it is then straightforward to show that: $\underline{\delta}_E^{MU(0)} > \underline{\delta}_E^{TT}$ and $\underline{\delta}_E^{MU(1)} < \underline{\delta}_E^{TT}$. This implies that MU equilibria exist both when TT exists and when TT does not exist.

3) Based on the proof of proposition 4, we know that TMU exists for $\delta_E \in [\underline{\delta}_E^{TMU}, \overline{\delta}_E^{TMU}]$. Using the definitions of $\overline{\delta}_E^{TMU}$and $\overline{\delta}_E^{TMU}$ simple calculations allow us to show that $\overline{\delta}_E^{TMU} \lesseqgtr \underline{\delta}_E^{TT}$, and the sign of this inequality may vary based on the values of γ, α and p. This implies that TMU equilibria *may* exist both when TT exists and when TT does not exist.

4) When $\delta_E < \underline{\delta}_E \equiv \min[\underline{\delta}_E^{MB}, \underline{\delta}_E^{MU(1)}, \underline{\delta}_E^{TMU}]$ no informative equilibria exist. To prove this, notice that given any strategy of U, for these values of δ_E the biased expert will always send $m_1 = 1$. Given this strategy of B, any putative equilibrium in which U is sending both signals, can never be an equilibrium since by the proofs of propositions 2 and 4, the DM will never hire after both messages, when the probability of misreporting of the B expert is too high. The only other plausible equilibrium involves U always sending $m_1 = 0$ and the DM hiring only after $m_1 = 0$. By the Proof of proposition 4, this can never be an equilibrium.