Good Lies*

Filippo Pavesi[†], Massimo Scotti[‡,§]

October 2019

Abstract

Decision makers often face uncertainty about the ability and the integrity of their advisors. If an expert is sufficiently concerned about establishing a reputation for being skilled and unbiased, she may truthfully report her private information about 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 she also learns about the expert's bias. Although truthtelling allows for better current decisions, it may lead to worse sorting outcomes. This occurs if misreporting is principally caused by biased experts driven by their conflict of interest rather than by unbiased experts attempting to signal their type. Whenever lying has these features, it may be welfare improving if the decision maker is sufficiently concerned about future choices.

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, Stephen Morris, Marco Ottaviani, Andrea Prat, and Toru Suzuki as well as seminar participants at University of Technology Sydney, University of Verona, University of Wollongong, Bocconi University (Dondena Workshop on Political Economy 2018), Erasmus University (Accounting & Economics Workshop 2017) THEMA at Cergy-Pontoise and conference participants at MWET 2017 in Lexington (KY), Lacea-Lames 2017 in Buenos Aires, Econometric Society Meeting 2017 in Algiers, GAMES 2016 in Maastricht, and the Game Theory Society Conference 2016 in Stony Brook (NY) for helpful comments and suggestions. All errors remain our own.

†School of Economics and Management, LIUC (Carlo Cattaneo University), C.so Matteotti, 22, 21053 Castellanza (VA), Italy; and Stevens Institute of Technology. School of Business, 1 Castle Point on Hudson, Hoboken, NJ 07030, USA. Email: fpvaesi@liuc.it

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

§Corresponding author

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*). From the politician's perspective, resolving this uncertainty can prove useful in deciding whether to replace the advisor or to continue to rely on her services. We investigate whether misreporting may have a positive effect on resolving this uncertainty, therefore possibly improving the quality of the politician's future decisions.

If the interaction between the politician and the advisor is based only on communication, and only implicit incentives are feasible, equilibrium behavior determines what the politician learns about the advisor's integrity and ability. For instance, if advisors are not sufficiently concerned about their reputation, biased experts will tend to ignore what their private information suggests and prescribe a policy that favors a tax break for a specific industry. In this case, observing such a recommendation may reveal information on the advisor's integrity. This may also suggest that in order to signal their integrity to the policy maker, unbiased experts may have an incentive to pander by refraining from recommending tax breaks even when their private information indicates that this could have positive effects on the economy.[1]

Lying thus reveals evidence about preferences which remains concealed if behavior is truthful. However, the overall quality of the expert's advice does not depend only on her integrity but also on her ability, and lying may limit learning about this latter dimension. Suppose, for example, that enacting a tax break suggested by the advisor actually leads to negative consequences for the economy. When advisors misreport, it is more difficult to attribute these errors to an actual lack of ability since, precisely because recommendations are not truthful, they reveal less information on the quality of the advisor's information.

Two natural questions thus arise in this setting. The first is whether facing advisors

[1]The idea that in the presence of costs of lying or reputational concerns, misreporting by both biased and unbiased experts may occur is well documented in the literature. Regarding misreporting by biased experts, Kartik (2009) shows how in a cheap talk game with conflicts of interest and costs of lying there is incomplete separation and pooling on higher messages. Morgan and Stocken (2003) find an analogous result proving that when there is uncertainty on the preferences of financial analysts, only bad information can be truthfully revealed. In terms of the empirical evidence, several papers document a tendency for financial analysts to bias their recommendations upwards that is consistent with the existence of conflicts of interests (Francis and Soffer, 1997; Lin and McNichols, 1998; Michaely and Womack, 1999).

Regarding misreporting by unbiased advisors, Morris (2001) finds that concerns for establishing a reputation for having policy preferences aligned with those of the decision maker, may lead experts to send politically correct evaluations. Likewise, in a principal-agent framework, Acemoglu et al. (2013) show how politicians pander by choosing populist policies that provide voters with a noisy signal that they are not right-wing politicians.

that sometimes lie in the current period, instead of honestly reporting their information, may allow politicians to make more informed decisions in the future. This boils down to understanding how the reporting strategies of the experts jointly affect learning about ability and integrity. The second question is whether unbiased advisors can improve upon the quality of the politician's future decisions by misreporting in order to signal their type.

To address these issues, we introduce a two period model that incorporates the key features of the example described above. Our analysis proves that some degree of misreporting may be preferable to truthful reporting because it leads to more informed future decisions. In particular, we find that lying by biased experts generally dominates lying by unbiased ones thereby unveiling that lies are "good" only if they are more likely to come from biased experts.

It is worth noting that the two-period model that we analyze is suggestive of cases in which reputational concerns become less effective in disciplining the behavior of an expert at some later stages of her interaction with the decision maker. In the discussion (Section 7.1) we illustrate how these features may also endogenously emerge when considering an infinite horizon.

The model we propose allows us to characterize situations in which the decision maker has concerns for learning about both the ability and the integrity of the expert she seeks advice from. The decision maker chooses a binary action in each period, and her payoff from the action depends on an unknown state of the world. Before choosing the action, 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 these two dimensions, 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.

In the second period, the expert will cease to have reputational concerns.[2] Hence, a biased expert will always send the recommendation that favors her conflict of interest. This makes it relevant to learn about the preferences of the expert that will be consulted in the second period. Moreover, the value of the expert's information also positively depends on her ability because a poorly informed unbiased expert is of little value for the decision maker. Therefore, the decision maker's posterior beliefs about integrity and ability jointly determine the expert's reputation which is essentially a measure of the value of the expert's advice in the second period. Accordingly, the value of reputation establishes whether the

[2] Throughout the paper we refer to reputational concerns and career concerns as synonyms.

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*), and we denote equilibria with this feature as *truthtelling equilibria* (TT). 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 her 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 informative 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. Hence, in these equilibria, while observing a certain recommendation reveals some information about integrity, learning about ability can never be sharper than in a truthtelling equilibrium 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. In particular, we find that this requires that the biased expert lies sufficiently more than the unbiased one. When this occurs, the decision maker learns about integrity without sacrificing too much learning about ability. 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.

More specifically, we first characterize a class of misreporting equilibria that always improve sorting with respect to truthtelling. This class, which we denote *Misreporting Biased* (MB), is defined as follows: 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; the decision maker retains the expert if and only if her recommendation is ex-post correct. These equilibria improve sorting with respect to truthtelling because misreporting by the biased expert leads her to make more mistakes, therefore increasing the likelihood that the expert that is retained after a correct recommendation is unbiased. At the same time, what the decision maker learns about the ability of an unbiased expert is the same as in TT, since in both classes of equilibria unbiased experts follow the same strategy of truthfully reporting their signals.

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. These equilibria, which we denote *Misreporting Unbiased* (MU), are of interest because they represent situations in which the unbiased advisor lies to signal her type to the decision maker. Indeed, like MB equilibria, this class also displays the feature that the expert's recommendation reveals some information about her integrity, and is characterized by the following strategies: the unbiased expert partially reveals her information about the state; the biased expert either truthtells or partially reveals her information depending on her level of career concerns; the decision maker retains the expert if and only if her recommendation is correct ex-post.

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 these MU equilibria would suggest a pattern of learning about ability and integrity, and hence a sorting effect that is similar to the one we have in MB equilibria. Nevertheless, 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, the sorting effect that comes from the unbiased expert's intention to signal her integrity is not sufficient to offset the sorting effect associated with truthtelling that derives from greater learning about ability. Indeed, in order for misreporting to have a positive effect on welfare, untruthful behavior must be predominantly associated with the biased expert. More specifically, we prove that the amount of lying of the biased expert must be strictly greater than that of the unbiased expert in order for MU equilibria to improve sorting.

Finally, in order to draw meaningful conclusions on welfare, we map the existence of equilibria with respect to a parameter that represents the relative weight that experts attribute to current versus future payoffs. While the existence of truthtelling requires a sufficiently high level of career concerns, we find that misreporting equilibria that can generate higher welfare than truthtelling exist only when career concerns are mild, and thus never coexist with truthtelling. Therefore, an increase in experts' career concerns may lead to a loss in welfare as we switch from equilibria with "good lies" to equilibria characterized by truthtelling, thereby indicating that it may not always be optimal for a decision maker to consult experts with high reputational concerns. In terms of policy implications, this result suggests that policies aimed at increasing career concerns may not necessarily always be desirable. On the contrary, reducing the value of experts' future payoffs may be welfare increasing, when the decision maker weights future decisions relatively more than current ones.

1.1 Literature Review

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; Bourjade and Jullien, 2011; Klein and Mylovanov, 2017; Schottmüller, 2019) or for preferences (Sobel, 1985; Benabou and Laroque,1992; Morris, 2001; Ely and Välimäki, 2003). A contribution of the present paper is to propose a model that incorporates both these sources of reputational concerns, rather than considering 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). In this respect, our work is related to Daley and Gervais (2017) that also consider a setting in which career concerns depend on two dimensions, namely skills and ethics. While we focus on comparing welfare across equilibria to determine how misreporting fares with respect to truthful behavior, Daley and Gervais (2017) analyze the properties of the unique inefficient equilibrium, that in their setting arises when ethics-related concerns are limited with respect to skills-related concerns.

Considering the literature about reputation for preferences, Morris (2001) and Ely and Välimäki (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 Välimäki (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 in a principal-agent model under the assumption that principals are short-run players. In fact, they show that if principals are long-run sufficiently-patient players, bad reputation is no longer a problem as principals can internalize the value of learning about the type of the agent and eventually achieve perfect screening. Our model also exploits this learning feature, but in a cheap-talk environment in which experts have no means to play payoff-relevant actions, and decision makers must learn about both the preferences and the ability of experts.

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, 2018). In particular, there is a parallel between our analysis and that of Frankel and Kartik (2018). 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 informative equilibria. Section 5 characterizes the misreporting equilibria that can potentially improve welfare relative to truthtelling. In Section 6, we discuss welfare implications and perform some comparative statics. In Section 7, we discuss the role of some key assumptions as well as the consequences of relaxing them. 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.[3] At the moment of choosing a_t, DM does not observe the realization of x_t. However, she can pay a fixed fee w_t and consult an expert who has access to a binary signal $s_t \in \{0, 1\}$ that is potentially informative about x_t. If the expert is hired, she observes s_t and then reports a message $m_t \in \{0, 1\}$ to DM. 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.[4]

We can think of DM's decision as an elected official's decision to enact a reform ($a_t = 1$) or to maintain the status quo ($a_t = 0$) when the outcome of the policy is uncertain, and we can think of the expert as a policy advisor. We assume that there is a finite pool of

[3]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.

[4]Analyzing optimal contracts is outside the scope of this paper, but it is worth noting that Aghion and Jackson (2016) consider optimal incentive mechanisms in a setting in which only dismissal can be used as a contractual device, a feature that is also present in our framework.

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.[5] 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-payoff equal to $w_t + a_t$, where also the component a_t is assumed to be relation-specific. Namely, a biased expert receives a_t in period t if and only if she has been hired by DM in period t.[6] 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 $R_t(a_t, x_t)$, that is, the payoff associated to her action gross of the fixed fee paid to the expert. 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

[5] 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.

[6] This is, for example, the case of a policy advisor that receives a compensation from a special interest group only if the politician favors that group while employing the advisor.

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 ex-ante expected value of Π_{DM}.

2.1 The Value Function and Reputational Concerns

We assume 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. As we will formally see in the next section, these beliefs determine the value of the incumbent's information in the second period. Intuitively, the more the incumbent expert is perceived to be unbiased and smart, the more useful her information is expected to be.

It is important to note that the assumption of a two-period horizon makes it relevant to learn about the integrity of the expert because a biased expert will always lie in her last period. In this respect, our two-period model is meant to capture those cases in which reputational concerns may be less effective in disciplining the behavior of biased experts at future dates. Section 7.1 discusses how these features may endogenously emerge when considering an infinite horizon.[7]

We denote the value of the incumbent's information in the second period with $V(m_1, x_1)$

[7] The two-period model also directly applies to settings in which the expert's reputation is relation specific, and ceases to be relevant at the end of the relation that is finite for exogenous reasons. This feature applies when the recommendations provided by experts, as well as the actual decisions, states and outcomes are difficult to observe or verify by external parties. For instance, this may be the case of an advisor that is hired at the beginning of a two-term mandate of a politician, and whose advisory record is unlikely to be publicly available at end of the politician's mandate. It also applies to corporate finance settings involving a firm's acquisition, in which sell-side advisors aim to establish a reputation for being valuable advisors in an initial phase, with the objective of obtaining an exclusive mandate for placing the firm on the market. After having obtained the mandate they hold the exclusive right to introduce prospective buyers to the firm, but once the owner follows the advisor's recommendation and the transaction takes place, reputation no longer plays a role since the selling owner no longer has a firm to sell, and naturally ceases to need the advisor's services. The seller's evaluation of the ability and integrity of the analyst remains part of the soft information that is relation specific, and is thus unlikely to be incorporated in prospective clients' beliefs about the analyst's qualities.

and refer to it as the value function. 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 $V(m_1, x_1) \geq V$. Second, we assume that the fixed fee that is paid to the expert at the beginning of the second period, w_2 is set equal to the value of the expert's information in the second period. Hence, if the incumbent is retained, she receives $w_2 = V(m_1, x_1)$. This is made for the sake of exposition and is without loss of generality. In Section 6, we discuss the implications of varying the share of $V(m_1, x_1)$ that the expert receives.[8]

All this implies that the incumbent will be concerned about being perceived as smart and unbiased for this maximizes $V(m_1, x_1)$, which in turn 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 described above. 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.[9]

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

[8]Assuming that w_2 depends positively on $V(m_1, x_1)$ is instrumental to generate reputational concerns that, conditional on the expert being retained, are continuously increasing in DM's belief that the expert is smart and unbiased.

[9]In terms of sorting, this structure makes the model qualitatively equivalent to the model with continuous action and quadratic loss function adopted, 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.

informative equilibrium, in each period, the expert's message must reveal some information about the state of the world.[10] This implies that in any informative equilibrium, m_t makes DM change her belief about x_t.[11] 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 $a_t(m_t) = m_t$.[12] With this in mind, we proceed by backward induction.

2.2.1 The Second Period

Reporting strategies and DM's action. An expert that is active in the second and last period does not have reputational concerns. 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.[13] 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 to the payoff that DM expects to attain in the second period thanks to the information of the incumbent.

At the beginning of the second period, DM observes (m_1, x_1) and updates her beliefs about the ability and the integrity of the incumbent. Since a biased expert always lies in the second period and a dumb expert always receives uninformative signals, only an expert that is both smart and unbiased adds value in the second period. Indeed, it is

[10]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.

[11]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.

[12]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.

[13]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.

straightforward to show that the equilibrium value of the incumbent's information in the second period reads:

$$V(m_1, x_1) \equiv E\left[R_2(a_2, x_2) \mid m_1, x_1\right] = \frac{r(2p-1)}{2} \Pr(U, S \mid m_1, x_1),$$

where $\Pr(U, S \mid m_1, x_1)$ is the joint probability that the incumbent expert is unbiased and smart conditional on observing m_1 and x_1, which can be interpreted as the reputation acquired by the incumbent expert at the end of the first period for being both unbiased and smart. In the reminder of the paper, we assume without loss of generality that $r = \frac{2}{(2p-1)}$ and focus on the case in which:

$$V(m_1, x_1) = \Pr(U, S \mid m_1, x_1).$$

In this way, the value of the incumbent's information in the second period coincides with her posterior reputation for being unbiased and smart, which in turn reflects what DM has learned about the the ability and the integrity of the incumbent after interacting with her in the first period.

For the sake of the subsequent analysis, it is convenient to write the conditional probability $\Pr(U, S \mid m_1, x_1)$ as follows:

$$\Pr(U, S \mid m_1, x_1) = \Pr(U \mid m_1, x_1) \Pr(S \mid U, m_1, x_1).$$

This allows us to work with the following expression of the value function:

$$V(m_1, x_1) = \widehat{\gamma}(m_1, x_1) \widehat{\alpha}(U, m_1, x_1), \tag{1}$$

where $\widehat{\gamma}(m_1, x_1) \equiv \Pr(U \mid m_1, x_1)$ and $\widehat{\alpha}(U, m_1, x_1) \equiv \Pr(S \mid U, m_1, x_1)$. Expression (1) shows that $V(m_1, x_1)$ can be decomposed in two components, one reflecting DM's posterior belief about the unbiased expert's ability and the other reflecting DM's posterior belief about the expert's integrity. As one would expect, $V(m_1, x_1)$ is strictly increasing in both these posterior beliefs.

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 and assuming $r = \frac{2}{(2p-1)}$, we obtain that the value of the information of a new expert reads:

$$V \equiv E\left[R_2(a_2, x_2)\right] = \gamma\alpha. \tag{2}$$

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.[14]

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:

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

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

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

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.

It turns out that a multitude of informative first-period reporting strategies are 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

[14]Note that because both $\widehat{V}(m_1, x_1)$ and V are strictly positive, DM always finds it optimal to consult an expert in period 2.

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, what we show to be true in Section 4.

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:[15]

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

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 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 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 ME equilibrium 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, it is useful to introduce two distinct measures of learning about $\widehat{\gamma}(m_1,x_1)$ and $\widehat{\alpha}(U,m_1,x_1)$, i.e., about the two components that determine $V(m_1,x_1)$.

Definition 1 *We let $|\widehat{\gamma}^{\sigma}(m_1,x_1)-\gamma|$ and $|\widehat{\alpha}^{\sigma}(U,m_1,x_1)-\alpha|$ respectively measure the amount of learning about integrity conditional on $(m_{1,}x_1)$ and the amount of learning about ability conditional on (U,m_1,x_1) in a putative equilibrium σ.*

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

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

Proposition 1 *(i)* $\left|\widehat{\gamma}^{ME}(m_1,x_1)-\gamma\right| > \left|\widehat{\gamma}^{TT}(m_1,x_1)-\gamma\right| = 0$ *for every* (m_1,x_1);

(ii) $\left|\widehat{\alpha}^{ME}(U,m_1,x_1)-\alpha\right| \leq \left|\widehat{\alpha}^{TT}(U,m_1,x_1)-\alpha\right|$*, with equality for every* (m_1,x_1) *if* U *truthfully reports her information, and strict inequality for at least one realization of* (m_1,x_1) *if* U *does not truthfully report her information.* {*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 weakly 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 advisor recommends a tax reform more often than an unbiased advisor, receiving a recommendation to cut taxes will rationally lead the politician to believe that the advisor is more likely to be biased than when she recommends maintaining the status quo.

As for ability, our relevant measure of learning is conditional on what the unbiased expert does in equilibrium. Note that in a truthtelling equilibrium, observing the recommendation of an unbiased expert 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 also the case in all those misreporting equilibria in which the unbiased expert tells the truth, implying that in these cases $\widehat{\alpha}^{ME}(U,m_1,x_1) = \widehat{\alpha}^{TT}(U,m_1,x_1)$. On the other hand, whenever we are in a misreporting equilibrium in which the unbiased expert misreports, at least some of her messages do not fully reflect her information. Hence, in this case inference about the ability of the unbiased 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 politician-advisor example, if a politician learns that her advisor is unbiased without learning enough about the unbiased advisor's ability, it is not obvious that the elected official will receive more informed policy advice in the future.

In Section 5, we show that there are 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. ■

4 Informative Equilibria

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 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.

4.1 Truthful Reporting by the Unbiased Expert

The following proposition characterizes this class of informative equilibria 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 reports her signals;

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

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

{*Proof in the Appendix*}.

To gather a better understanding of how each of these equilibria arises, first consider TT. We already know by Proposition 1 that in a truthtelling equilibrium messages are not informative about integrity, that is:

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

Concerning beliefs $\widehat{\alpha}^{TT}(U, m_1, x_1)$, it is straightforward to show that:

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

where $\underline{\alpha} = \frac{\alpha(1-p)}{1-q}$ and $\overline{\alpha} = \frac{\alpha p}{q}$. Hence, correct (incorrect) messages cause DM's belief about ability to increase above (decrease below) the prior α.

It is then immediate to verify that:

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

where $\underline{V} = \frac{\gamma\alpha(1-p)}{1-q}$ and $\overline{V} = \frac{\gamma\alpha p}{q}$.

Expression (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 if she is more concerned about her continuation payoff than her current payoff, that is, if δ_E is sufficiently large. 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.

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 uncertainty about the characteristics of the expert (i.e., ability) that creates the right incentives to fully reveal information about the state of the world.[16]

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 $s_1 = 0$. However, as we formally show in the appendix, there always exists a scalar $\underline{\delta}_E^{MB} \in \left(0, \underline{\delta}_E^{TT}\right)$ such that if $\underline{\delta}_E^{MB} \leq \delta_E < \underline{\delta}_E^{TT}$ (i.e., for intermediate values of

[16]Section 7.2 explores this issue in further detail.

career concerns) the MB equilibria described in proposition 2 can be supported.

To understand how MB equilibria arise, let us note that, given the reporting strategies of biased and unbiased experts in MB, observing $m_1 = 0$ ($m_1 = 1$) signals that the expert is likely to be unbiased (biased). In particular, it is easy to verify that the beliefs about integrity satisfy:

$$\widehat{\gamma}^{MB}(1,0) < \widehat{\gamma}^{MB}(1,1) < \gamma < \widehat{\gamma}^{MB}(0,1) = \widehat{\gamma}(0,0)^{MB}. \tag{8}$$

Since U truthtells in MB, beliefs $\widehat{\alpha}^{MB}(U, m_1, x_1)$ follow the same pattern as in TT:

$$\underline{\alpha} = \widehat{\alpha}^{MB}(U,0,1) = \widehat{\alpha}^{MB}(U,1,0) < \alpha < \widehat{\alpha}^{MB}(U,1,1) = \widehat{\alpha}^{MB}(U,0,0) = \overline{\alpha}. \tag{9}$$

The chains of inequalities described by (8) and (9) immediately imply that $V^{MB}(1,0) < V < V^{MB}(0,0)$, leading the decision maker to retain an incumbent that provides a correct evaluation when sending $m_1 = 0$, and replacing one that provides an incorrect one when reporting $m_1 = 1$. However, establishing whether $V^{MB}(0,1)$ and $V^{MB}(1,1)$ are greater than or less than V depends on the degree of misreporting. Indeed, as $\lambda_{B,0}$ decreases and B lies more frequently, the power of message $m_1 = 0$ ($m_1 = 1$) to signal that the expert is likely to be unbiased (biased) increases, and hence $\widehat{\gamma}^{MB}(0, x_1)$ increases ($\widehat{\gamma}^{MB}(1, x_1)$ decreases). At the same time, $\widehat{\alpha}^{MB}(U, m_1, x_1)$ does not change since it depends only on the reporting strategy of U. As long as $\lambda_{B,0}$ is not too small (i.e., if B does not lie too frequently), both $\widehat{\gamma}^{MB}(1,1)$ and $\widehat{\gamma}^{MB}(0,1)$ remain close enough to the prior γ and the following inequalities hold true:

$$V^{MB}(1,0) < V^{MB}(0,1) < V < V^{MB}(1,1) < V^{MB}(0,0). \tag{10}$$

When the above inequalities are satisfied, DM finds it optimal to retain the incumbent after a correct message and replace her after an incorrect one, exactly as in TT equilibria.

It is worth stressing that this is the only possible strategy of the decision maker that is consistent with MB equilibria. Indeed, if both $\widehat{\gamma}^{MB}(1,1)$ and $\widehat{\gamma}^{MB}(0,1)$ moved farther away from the prior γ as result of $\lambda_{B,0}$ becoming too small (i.e., as a result of B lying too frequently), we would have that $V^{MB}(1,1) < V < V^{MB}(0,1)$, and it would never be optimal for DM to retain the incumbent after receiving $m_1 = 1$.[17] Thus, the unbiased expert would never send $m_1 = 1$, therefore deviating from her prescribed truthtelling strategy.

[17]Notice that the result that DM's strategy is unique relies on the observation that $V^{MB}(1,1)$ and $V^{MB}(0,1)$ can never both be above or below V at the same time. This follows from the proof of proposition 2 in the appendix, in which we show that there exists a unique value of $\lambda_{B,0}$ for which both expressions are equal to V.

Finally, to see that B retains an incentive to sometimes truthfully report signal $s_1 = 0$ despite her career concerns being smaller than in TT (i.e., for $\delta_E < \underline{\delta}_E^{TT}$), note that (8) and (9) imply that $V^{MB}(0,0) > V^{TT}(0,0) = V^{TT}(1,1) > V^{MB}(1,1) > V$. Hence, since DM retains the incumbent expert after a correct message and replaces her after an incorrect one, the biased expert obtains a larger reward for truthfully reporting $s_1 = 0$ in MB than in TT.

4.2 Misreporting by the Unbiased Expert

We now focus on the class of 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 this class of informative equilibria.

Proposition 3 *Each equilibrium in which U misreports, which we denote Misreporting Unbiased (MU), has the following properties:*

i) U partially reveals one signal and truthfully reports the other signal; ii) B always truthfully reports $s_1 = 1$ and reports $s_1 = 0$ with probability $\lambda_{B,0} \in (0,1]$; in all cases, B reports the message that is falsely reported by U less often than U does; iii) DM retains the incumbent if and only if $m_1 = x_1$. {*Proof in the Appendix*}

Notice that misreporting a signal implies sending a message that is likely to be incorrect ex-post. Hence, lying is likely to negatively affect DM's beliefs about the ability of the unbiased expert and thus the expert's expected payoff. Therefore, the only reason for the unbiased expert to lie is that 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, therefore compensating for the loss in reputation for ability.

Note that we can group equilibria in which U misreports into two broad classes: Equilibria in which U partially reveals $s_1 = 1$ and truthfully reveals $s_1 = 0$; and those in which U truthfully reveals $s_1 = 1$ and partially reveals $s_1 = 0$.[18] For an intuition of the underlying

[18]Notice that MU equilibria in which U partially reveals $s_1 = 0$ and truthfully reveals $s_1 = 1$ are characterized by the fact that observing $m_1 = 1$ results in a positive update on integrity. To support these equilibria, B's strategy must be such that if she misreports $s_1 = 0$, she must do so with lower 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. Since 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.

forces at play, let us focus on the former equilibria.[19]

For the sake of exposition, let us consider the case in which B truthtells (i.e., $\lambda_{B,0} = 1$) and let us denote with $\lambda_{U,1}$ the probability with which U partially reports $s_1 = 1$. In these equilibria, message $m_1 = 0$ ($m_1 = 1$) signals that the incumbent is likely to be unbiased. In particular, it can be verified that:

$$\widehat{\gamma}^{MU}(1,1) = \widehat{\gamma}^{MU}(1,0) < \gamma < \widehat{\gamma}^{MU}(0,0) < \widehat{\gamma}^{MU}(0,1). \tag{11}$$

Note that as $\lambda_{U,1}$ decreases and U lies more frequently, the power of message $m_1 = 0$ ($m_1 = 1$) to signal that the expert is likely to be unbiased (biased) increases, and hence $\widehat{\gamma}^{MU}(0, x_1)$ increases ($\widehat{\gamma}^{MU}(1, x_1)$ decreases). At the same time, one can show that beliefs $\widehat{\alpha}^{MU}(U, m_1, x_1)$ satisfy the following pattern:

$$\underline{\alpha} = \widehat{\alpha}^{MU}(U,1,0) < \widehat{\alpha}^{MU}(U,0,1) < \alpha < \widehat{\alpha}^{MU}(U,0,0) < \widehat{\alpha}^{MU}(U,1,1) = \overline{\alpha}. \tag{12}$$

Hence, as one would expect, $\widehat{\alpha}^{MU}(U, m_1, x_1)$ increases above its prior α when the message is correct, and decreases below α when the message is incorrect. However, the inference that DM makes about the ability of U conditional on observing $m_1 = 0$ is now hindered by the fact that U sometimes falsely reports message $m_1 = 0$, which explains why $\widehat{\alpha}^{MU}(U,0,0) < \widehat{\alpha}^{MU}(U,1,1)$ and $\widehat{\alpha}^{MU}(U,1,0) < \widehat{\alpha}^{MU}(U,0,1)$. In particular, as $\lambda_{U,1}$ decreases and U lies more frequently, both $\widehat{\alpha}^{MU}(U,0,1)$ and $\widehat{\alpha}^{MU}(U,0,0)$ tend to α. Conversely, as $\lambda_{U,1}$ increases and U lies less frequently, $\widehat{\alpha}^{MU}(U,0,1)$ and $\widehat{\alpha}^{MU}(U,0,0)$ tend to $\underline{\alpha}$ and $\overline{\alpha}$ respectively.

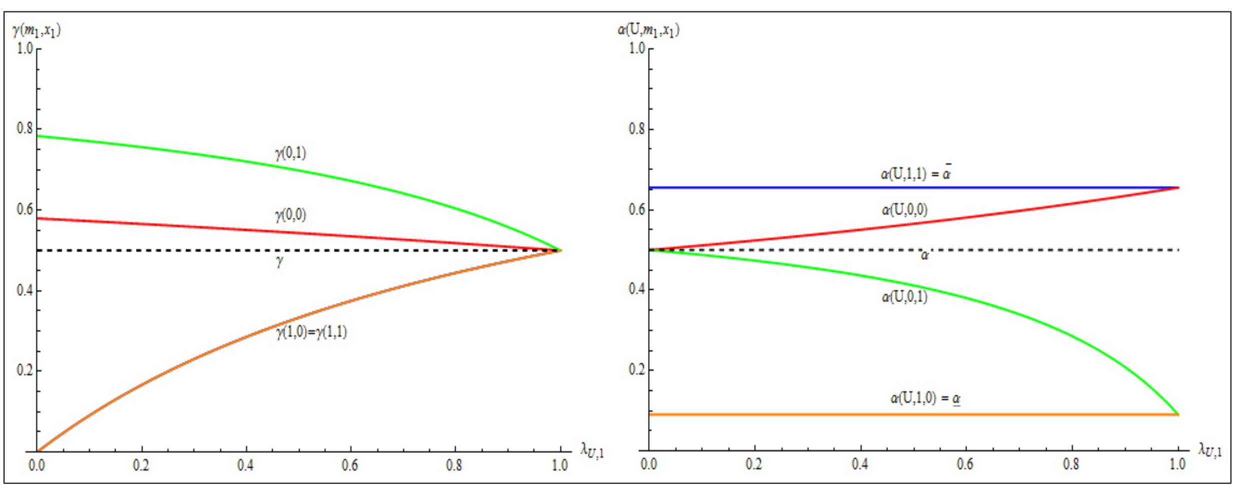

Figure 1: Beliefs $\gamma(m_1, x_1)$ and $\alpha(U, m_1, x_1)$ as a function of U's probability of truthfully reporting signal $s_1 = 1$.

[19] A similar reasoning applies when considering MU equilibria in which U truthfully reveals $s_1 = 1$ and partially reveals $s_1 = 0$.

Figure 1 provides a visual representation of how $\widehat{\gamma}^{MU}(m_1, x_1)$ and $\widehat{\alpha}^{MU}(U, m_1, x_1)$ vary with $\lambda_{U,1}$.

While (11) and (12) immediately imply that $V^{MU}(1,0) < V < V^{MU}(0,0)$, the discussion above serves to illustrate that when $\lambda_{U,1}$ is sufficiently high, it is also true that $V^{MU}(0,1) < V < V^{MU}(1,1)$. Hence, when $\lambda_{U,1}$ is sufficiently high, DM finds it optimal to retain the incumbent after a correct message and replace her after an incorrect one.[20]

Finally, we note that in order to induce the biased expert to report her signals truthfully, career concerns represented by δ_E must be sufficiently high.

5 Can Misreporting Be Welfare Improving?

We are now ready to address our key question. Can equilibria with misreporting deliver a higher welfare than equilibria with truthtelling? We know from the analysis of Section 3 that equilibria with some misreporting can generate better sorting and therefore even higher welfare if the decision maker is particularly concerned about future decisions. We now show that, relative to TT, misreporting can have positive effects on sorting and thus can potentially improve welfare whenever the biased expert lies more than the unbiased expert. This is obviously the case of MB equilibria, and also the case of some MU equilibria.

5.1 Misreporting Biased vs Truthtelling Equilibria

We begin by comparing how MB equilibria fare in terms of welfare with respect to TT equilibria. A direct implication of Proposition 1 is that there is strictly more learning in MB with respect to TT. Indeed, since MB equilibria involve misreporting only on behalf of biased experts, misreporting leads to more learning about integrity and the same amount of learning about ability with respect to truthtelling. The following proposition is a direct consequence of this observation.

Proposition 4 *MB always improves sorting with respect to TT. {Proof in the Appendix}.*

To gather further intuition for this result, 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

[20] As we formally show in the proof of Proposition 3, this is the only DM strategy that is consistent with an equilibrium in which U lies. The logic of the proof follows a similar arguement as the one illustrated in Section 4.1 to describe the existence of MB equilibria.

the following expressions representing the ex-ante second-period expected payoffs in MB and TT respectively:

$$\begin{aligned} E_0^{MB}\left[R_2(a_2,x_2)\right] &= \Pr(1,1|MB)V^{MB}(1,1)+\Pr(0,0|MB)V^{MB}(0,0) \\ &\quad +\Pr(0,1|MB)V+\Pr(0,1|MB)V, \end{aligned} \tag{13}$$

$$\begin{aligned} E_0^{TT}\left[R_2(a_2,x_2)\right] &= \Pr(1,1|TT)\overline{V}+\Pr(0,0|TT)\overline{V} \\ &\quad +\Pr(1,0|TT)V+\Pr(0,1|TT)V. \end{aligned} \tag{14}$$

Proposition 4 states that $E_0^{MB}\left[R_2(a_2,x_2)\right]-E_0^{TT}\left[R_2(a_2,x_2)\right]$ is always strictly positive. Note that this difference can be decomposed into two components. First, consider the difference between the bites of (13) and (14) 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:

$$\begin{aligned} &\Pr(1,0|MB)V+\Pr(0,1|MB)V-\left[\Pr(1,0|TT)V+\Pr(0,1|TT)V\right]= \\ &=\left[\Pr(m_1\neq x_1,B|MB)-\Pr(m_1\neq x_1,B|TT)\right]V>0. \end{aligned} \tag{15}$$

Expression (15) highlights that the replacement component is always positive. This occurs because the probability of replacing an unbiased expert is the same in both equilibria because the unbiased expert follows the same strategy in both equilibria; while the probability of correctly replacing a biased expert is strictly higher in MB than in TT, because in MB the biased expert misreports with positive probability and hence her chances of making a mistake are larger than in TT.

Now, consider the difference between the bites of (13) and (14) 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:

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

Now recall that $V^{\sigma}(m_1,x_1)=\widehat{\gamma}^{\sigma}(m_1,x_1)\widehat{\alpha}^{\sigma}(U,m_1,x_1)$, and note that $\widehat{\gamma}^{\sigma}(m_1,x_1)=\frac{\gamma\Pr(m_1,x_1|U,\sigma)}{\Pr(m_1,x_1|\sigma)}$ and $\widehat{\alpha}^{\sigma}(U,1,1)=\widehat{\alpha}^{\sigma}(U,0,0)=\overline{\alpha}$ for $\sigma\in\{TT,MB\}$. Hence, the continuation component can equivalently be written as follows:

$$\begin{aligned} &\left[\Pr(1,1|U,MB)+\Pr(0,0|U,MB)\right]\gamma\overline{\alpha}+ \\ &-\left[\Pr(1,1|U,TT)+\Pr(0,0|U,TT)\right]\gamma\overline{\alpha}=0, \end{aligned} \tag{16}$$

where the equality follows from the fact that, since U truthtells in MB, $\Pr(m_1, x_1|U, MB) = \Pr(m_1, x_1|U, TT)$.

Therefore, the fact that there is altogether more learning in MB with respect to TT implies that misreporting increases the chances of correctly replacing biased experts with respect to truthtelling (replacement component), without affecting the expected value of the information of incumbents that are retained (continuation component).

Thus, Proposition 4 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 leads to 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.

5.2 Misreporting Unbiased vs Truthtelling Equilibria

A natural question is whether also equilibria in which an unbiased expert lies have the potential to improve sorting and hence the expected utility of DM relative to TT. From Proposition 1, we know that unlike MB equilibria, MU equilibria are not clearly superior to TT in terms learning since misreporting by the unbiased expert leads to more learning on integrity and less on ability. In particular, we now show that while misreporting has a positive effect on the chances of correctly replacing a biased expert, lying by the unbiased expert leads to a drop in the chances of correctly retaining an unbiased expert of high ability in the second period. The following proposition states that for the former positive effect to prevail over the latter negative one, U must lie strictly less than B. [21]

Proposition 5 *The only MU equilibria that can improve sorting with respect to TT (that we denote MU^*) are characterized by: (i) B truthfully reporting $s_1 = 0$ with probability $\lambda_{B,0} \in (0, 1)$; and (ii) U truthfully reporting $s_1 = 1$ with probability $\lambda_{U,1} \in (\lambda_{B,0}, 1)$. {Proof in the Appendix}.*

In order to gather a deeper intuition for this result, we can break up the net welfare gain of MU 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):

$$\begin{gathered}[\Pr(m_1 \neq x_1 \mid MU) - \Pr(m_1 \neq x_1 \mid TT)]\, V = \\ = \tfrac{1}{2}\gamma\alpha(2q-1)[(1-\gamma)\,(1-\lambda_{B,0}) - \gamma(1-\lambda_{U,1})],\end{gathered} \tag{17}$$

[21] As explained intuitively in footnote (18) and shown formally in the appendix, MU equilibria in which U misreports $s_1 = 0$ and truthfully reports $s_1 = 1$ necessarily require that U lies strictly more than B. Therefore, these MU equilibrium can never improve sorting over truthtelling.

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

$$\begin{aligned}&\Pr(1,1 \mid U, MU)\gamma\overline{\alpha} + \Pr(0,0 \mid U, MU)\gamma\alpha_{Low} + \\ &- \Pr(1,1 \mid U, TT)\gamma\overline{\alpha} - \Pr(0,0 \mid U, TT)\gamma\overline{\alpha} = \\ &= \tfrac{1}{2}\gamma[(1-\lambda_{U,1})((1-q)\alpha_{Low} - q\overline{\alpha}) + q(\alpha_{Low} - \overline{\alpha})]\end{aligned} \tag{18}$$

where $\alpha_{Low} \equiv \alpha^{MU}(U,0,0) < \overline{\alpha}$.

First, notice that the replacement component (17) is decreasing in $\lambda_{B,0}$ and increasing in $\lambda_{U,1}$. Importantly, it is strictly positive if $\lambda_{U,1}$ is sufficiently higher than $\lambda_{B,0}$. The intuition is straightforward. While replacing a biased expert is beneficial to the decision maker, replacing an unbiased one is costly. The two bites $(1-\gamma)(1-\lambda_{B,0})$ and $\gamma(1-\lambda_{U,1})$ in (17) represent the probabilities of lying by a biased and an unbiased expert respectively. So, these are also the probabilities with which the two types of experts make mistakes and are replaced. Therefore, as long as $\lambda_{U,1}$ is high enough with respect to $\lambda_{B,0}$, the replacement component (17) becomes positive because it is more likely that the decision maker replaces a biased expert rather than an unbiased one.

Second, note that the continuation component (18) is always negative. This is easily seen by observing that $\alpha_{Low} < \overline{\alpha}$ and $\frac{1}{2} < q < 1$. There are two reasons why this is the case. The first reason is that, relative to TT, there is a drop in the chances of retaining an unbiased expert of high ability $\overline{\alpha}$. Namely, $\Pr(1,1 \mid U, MU) < \Pr(1,1 \mid U, TT)$. This is due to the fact that the misreporting strategy of the unbiased expert implies that realization $(m_1, x_1) = (1,1)$ occurs less often in MU than in TT. The second reason is that realization $(m_1, x_1) = (0,0)$ occurs more often in MU but generates a posterior α_{Low} which is lower that the posterior $\overline{\alpha}$ that arises in TT.

Finally, notice that as $\lambda_{U,1}$ increases while $\lambda_{B,0} < 1$, the replacement component becomes strictly positive while the negative continuation component shrinks to zero because α_{Low} approaches $\overline{\alpha}$. For an intuition of this latter effect, simply observe that as $\lambda_{U,1}$ tends to 1, an MU equilibrium tends towards an MB equilibrium, and we know from section 5.1 that for MB equilibria the continuation component is equal to zero. This implies that as $\lambda_{U,1}$ increases, there is a point at which the positive sorting due to the replacement component dominates the negative one resulting from the continuation component.

6 Welfare Implications and Comparative Statics

Section 5 suggests that misreporting equilibria have the potential to improve welfare so long as the biased expert lies sufficiently more than the unbiased one. Hence, lies are potentially good only if they are more likely to come from a biased expert. So, equilibria in which the biased expert sometimes lies and the unbiased expert always truthtells (i.e., MB equilibria) always improve sorting and thus have the potential to improve welfare relative to truthtelling. On the other hand, attempts by the unbiased expert to signal her own type through lying are likely to be costly in terms of welfare since only a sub-set of MU equilibria (i.e., MU^* equilibria) does better than TT in terms of sorting.

The next proposition establishes that MB and MU^* never coexist with TT.

Proposition 6 *There exist $\underline{\delta}_E, \overline{\delta}_E \in (0,1)$ with $\underline{\delta}_E < \overline{\delta}_E$ such that:*

a) For $\delta_E \in [\underline{\delta}_E, \overline{\delta}_E)$, there always exists a non-empty set of informative equilibria that includes at least MB or MU^* but does not include TT;

b) For $\delta_E \in [\overline{\delta}_E, 1]$, there always exists a non-empty set of informative equilibria that includes TT but does not include MB and MU^*.

{*Proof in the Appendix*}.

The intuition of proposition 6 goes as follows. We know from Section 4 that MB equilibria exist only for $\underline{\delta}_E^{MB} < \delta_E < \underline{\delta}_E^{TT}$ and thus never coexist with TT. Instead, MU equilibria may coexist with TT since they require δ_E to be sufficiently high to induce the biased expert to reveal her information (either fully or partially depending on the strength of career concerns). Now, let us focus on those MU equilibria in which the unbiased expert truthfully reports $s_1 = 1$ with probability strictly less than 1.[22] Intuitively, in these equilibria, when $\delta_E > \underline{\delta}_E^{TT}$, a biased expert has an incentive to truthtell that is even stronger than the incentive she would have in a TT equilibrium. This is so because now message $m_1 = 0$ (i.e., the message the biased expert likes less) provides a signal of integrity and becomes attractive. Hence, when $\delta_E > \underline{\delta}_E^{TT}$, this class of MU equilibria is characterized by the biased expert truthtelling.

Proposition 6 thus illustrates that 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. In particular, the following general welfare results apply. When δ_E is sufficiently high, TT exists and is welfare maximizing with respect to the other informative equilibria that exist

[22] We know that those MU equilibria in which the unbiased expert truthfully reports $s_1 = 0$ with probability strictly less than 1 never improve sorting.

for these high values of career concerns. However, when experts do not care enough about future payoffs, truthtelling breaks down but there always exist other equilibria that involve some degree of misreporting (i.e., either MB or MU^*) that may generate higher levels of sorting with respect to truthtelling.

A further implication of Proposition 6 is that when the decision maker puts enough weight on future decisions, an increase in experts' career concerns may lead to a loss in welfare as we switch from equilibria with "good lies" to equilibria characterized by truthtelling. This interpretation provides some insight in assessing the validity of policies aimed at increasing experts' career concerns and consequently their incentives to truthfully report their information.

These policies may for example be implemented by loosening regulatory pressure relative to the fees that experts require for their services. Rather counter-intuitively, such interventions may actually backfire and, if the decision maker is sufficiently concerned about the future, even lead to a welfare loss. Indeed, increasing career concerns by allowing experts to obtain a greater share of the expected value of their information may reduce the incentives to misreport, ultimately implying that equilibria characterized by "good lies" such as MB and MU^* may cease to exist.

To see this point, we can slightly modify our model by assuming that the fixed fee that the expert receives in the second period is a share $\rho \in (0, 1]$ of the expected value of her information, i.e., $w_2 = \rho V(m_1, x_1)$. In this case, the truthtelling condition that determines $\overline{\delta}_E$ now reads:

$$(1-\delta_E)\left(\overline{w}+a(0)\right)+\delta_E q\left(\rho\overline{V}+1\right) \geq (1-\delta_E)\left(\overline{w}+a(1)\right)+\delta_E(1-q)\left(\rho\overline{V}+1\right),$$

which simplifies to:

$$\delta_E \geq \overline{\delta}_E = \frac{1}{\rho\left(2q-1\right)\overline{V}+2q}.$$

The expression above illustrates that a policy aimed at increasing ρ leads to a decrease in $\overline{\delta}_E$, therefore expanding the parameter space for which potentially welfare-inferior truthtelling equilibria exist.

7 Discussion

In this section, we discuss some of the key features and assumptions of the model. We begin by providing an intuition of how our results from the two period model could be extended to an infinite horizon setting. We then illustrate the role that uncertainty on ability has in providing incentives for truthtelling that are not present when the decision maker faces uncertainty exclusively on the expert's integrity. Finally, we consider how our results would vary if experts and decision makers were to have the same discount factor, and conclude by analyzing the case in which the signal space is continuous rather than binary.

7.1 Considering a Longer Horizon

In our two-period model, a biased expert always lies in the last period. This makes learning about the preferences of the incumbent expert in the first period (and the related sorting between biased and unbiased ones) particularly relevant. One may suspect that with a longer horizon, career concerns may discipline the behavior of experts and support an equilibrium in which all types of experts have an incentive to always truthfully report their information. In this case, any learning about preferences would become irrelevant, and there would be no need to sort biased from unbiased experts since both behave efficiently.

In Section A.9 of the Appendix, we formally prove that there cannot exist Markovian equilibria in which both biased and unbiased experts truthtell in every period of an infinite horizon game.[23] For an intuition, note that in a putative truthtelling equilibrium there would be no learning about the preferences of an incumbent expert while there would be learning about her ability. Now suppose that the incumbent makes a string of successful calls that brings her reputation for being smart (and hence her overall reputation for being smart and unbiased) well above the prior (i.e., well above the starting reputation of a challenger). At this point, an incorrect recommendation would lower the posterior reputation of the incumbent but not enough to fall below the prior. Hence, after such a history, which has positive probability of occurring on the equilibrium path, the expert would be retained even after providing a wrong recommendation. A biased expert would therefore have an incentive to deviate from her truthful strategy and send $m = 1$ irrespective of her signal.

The reasoning above suggests that if informative Markovian equilibria exist, they are characterized by either the biased expert or the unbiased one lying. Since the biased expert has strictly greater incentives to lie than the unbiased one, we should expect that there exist equilibria in which the biased expert will lie with higher probability than the unbiased

[23] Notice that when considering non-Markovian sequential equilibria, there may exist equilibria in which a sufficiently patient decision maker may achieve approximately his full information value as in Ely and Välimäki (2003). These equilibria are very demanding in terms of information requirements (and therefore rather implausible) because unlike Markovian equilibria, they require that the decision maker knows the full history of recommendations made by the expert and not just her reputation. Aghion and Jackson (2016) adopt a similar line of reasoning to rule out trigger-strategy equilibria in the infinite horizon version of their game without commitment.

one. This provides a justification for the fact that the decision maker may benefit from sorting biased experts from unbiased ones. Intuitively, sorting is all the more effective, the more the biased expert lies relative to the unbiased one thereby creating a learning effect that is equivalent to the one studied in the two-period model.

7.2 The Role of Reputation for Ability

It is worth noting that informative equilibria would not exist if reputational concerns were only related to preferences. It is the fact that reputational concerns encompass two dimensions 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. 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, indeed, that our setup is equivalent to assuming that U's reputational concerns are maximum, meaning that the unbiased expert is not concerned at all about current decisions. When we set $\alpha = 1$, as prescribed by Morris (2001), 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.

7.3 Different δ for the Expert and the Decision Maker

Throughout the analysis we assumed that the expert and the decision maker could assign different weights to current versus future payoffs. A natural question is therefore to ask how our results would be affected by imposing that all agents share the same δ. For example, a high value of δ represents the case in which DM assigns relatively more weight to future versus present decisions. We know that this implies that DM would prefer her advisor to play an equilibrium with a strong sorting effect, i.e. either MB or MU^*. However, as shown in Proposition 6, a high value of δ means greater career concerns for the expert, and hence the fact that MB or MU^* may not be sustainable in equilibrium while TT may indeed exist. Paradoxically, truthtelling would be played when it is least needed. In this case, a policy aimed at reducing regulatory pressure on expert fees, such as the one mentioned in section 6, may produce the desired result of going from a truthtelling equilibrium to one in which there is some misreporting.

7.4 Continuous Signals

While we presented our results assuming a binary signal structure, the main features of our results continue to hold even if the signal structure is continuous and the action space is binary. As shown in Garfagnini et al. (2014), in such a setting we can restrict attention to a binary message space, where informative equilibria are now characterized by a threshold on the signal space such that for realizations above (below) this critical value the expert will send positive (negative) recommendations. In this context, it is straightforward to show that the qualitative features of Proposition 6 continue to hold. Indeed, for values of delta above a critical threshold there exists a truthful equilibrium which involves both biased and unbiased experts sending positive recommendations, only for those signal realizations for which they truly believe that choosing $a_1 = 1$ will maximize $E\left[R_1(a_1, x_1)\right]$, and sending a negative recommendation otherwise. As in the binary signal space model, for intermediate values of δ_E, these truthful equilibria no longer exist since the biased expert will prefer to send a positive recommendation even for lower signal realizations. This occurs because the biased expert is willing to trade-off the lower chances of providing a correct recommendation, and therefore being rehired, with the immediate benefit of inducing her preferred action. Thus, for these values of δ_E, it is always possible to construct an equilibrium that resembles MB in which the unbiased expert continues to behave truthfully (i.e., by using the same threshold as in the truthful equilibrium) while the biased expert sets a strictly lower threshold. As in our binary signal setting, these MB equilibria involve greater sorting with respect to TT, and therefore can be welfare improving with respect to truthful reporting if the DM is sufficiently concerned about the future.[24]

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

[24]A more formal analysis of this result is provided in Sction A.10 of the Appendix.

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 may suffice to support truthful reporting by an expert in the early stages of her interaction with the decision maker. However, we point out that this outcome may not necessarily be the one 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., her bias). In particular, with respect to truthtelling, misreporting equilibria lead to more learning about integrity and possibly 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] Acemoglu, D., Egorov, G., and Sonin, K. 2013. A Political Theory of Populism. *Quarterly Journal of Economics*, 128 (2), 771-805.

[2] Aghion, P., Jackson, M.O. 2016. Inducing Leaders to Take Risky Decisions: Dismissal, Tenure, and Term Limits. *American Economic Journal: Microeconomics*, 8(3), 1-38.

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

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

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

[6] Bourjade, S., Jullien, B. 2011. The Role of Reputation and Transparency on the Behavior of Biased Experts. *The RAND Journal of Economics,* 42, 575-594.

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

[8] Daley, B., Gervais, S. 2017. The Equilibrium Value of Employee Ethics, *Working Paper*.

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

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

[11] Francis, J., Soffer, L. 1997. The Relative Informativeness of Analysts' Stock Recommendations and Earnings Forecast Revisions. *Journal of Accounting Research*, 35, 193-211.

[12] Frankel, A., Kartik, N. 2018. Muddled Information. *Journal of Political Economy (forthcoming)*.

[13] Garfagnini, U., Ottaviani, M., and Sorensen, P.N. 2014. Accept or Reject? An Organizational Persepective. *International Journal of Industrial Organization*, 34, 66-74.

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

[15] Kartik, N. 2009. Strategic Communication with Lying Costs. *Review of Economic Studies*, 76 (4), 1359-1395.

[16] Klein M., Mylovanov T. 2017. Will Truth Out? An Advisor's Quest to Appear Competent. *Journal of Mathematical Economics,* 72, 112-121.

[17] 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.

[18] Lin, H.W., McNichols, M.F. 1998. Underwriting Relationships, Analysts' Earnings Forecasts and Investment Recommendations. *Journal of Accounting and Economics*, 25, 101-127.

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

[20] Morgan, J., Stocken, P.C. 2003. An Analysis of Stock Recommendations. *The RAND Journal of Economics*, 34(1), 183-203.

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

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

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

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

[25] Schottmüller, C. 2019. Too Good to be Truthful: Why Competent Advisors are Fired. *Journal of Economic Theory,* 181, 333–360.

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

[27] 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 signal s_1 truthfully. That is, $\lambda_{i,s_1} = \Pr(m_1 = s_1 \mid s_1, i)$.

c) The expression *expert i misreports signal s_1* denotes the case in which $\lambda_{i,s_1} < 1$;

d) The expression *expert i partially reports signal s_1* denotes the case in which $0 < \lambda_{i,s_1} < 1$;

e) The expression *expert i truthfully reports signal s_1* denotes the case in which $\lambda_{i,s_1} = 1$.

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

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 continuation equilibrium of the second period: 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 the most informative continuation equilibrium in which the unbiased expert acts in the interest of the DM and truthfully reveals her signal. ■

Let $V \equiv E\left[R_2(a_2, x_2)\right]$ and $V(m_1, x_1) \equiv E\left[R_2(a_2, x_2) \mid m_1, x_1\right]$.

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:

$$V = \gamma\alpha,$$

$$V(m_1, x_1) = \widehat{\gamma}(m_1, x_1)\widehat{\alpha}(U, m_1, x_1).$$

Note that 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 retains the incumbent whenever $V(m_1, x_1) \geq V$ and fires her otherwise. ■

It is immediate to verify that $V(0, x_1)$ and $V(1, x_1)$ are respectively strictly decreasing and strictly increasing in $\lambda_{B,0}$ for any $x_1 = 0, 1$; and that $V(0, 1)$ and $V(1, 1)$ are respectively strictly decreasing and strictly increasing in $\lambda_{U,1}$.

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 $\left[V(m_1, x_1) + 1\right]\imath(m_1, x_1)$, 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:

$$\delta_E \pi_{2,B}(0,0) - (1-\delta_E) - \delta_E \pi_{2,B}(1,0) \geq 0 \text{ if } s_1 = 0, \quad (19)$$

$$(1-\delta_E) + \delta_E \pi_{2,B}(1,1) - \delta_E \pi_{2,B}(0,1) \geq 0 \text{ if } s_1 = 1, \quad (20)$$

For an unbiased expert instead, we have:

$$\pi_{2,U}(0,0) - \pi_{2,U}(1,0) \geq 0 \text{ if } s_1 = 0, \quad (21)$$

$$\pi_{2,U}(1,1) - \pi_{2,U}(0,1) \geq 0 \text{ if } s_1 = 1, \quad (22)$$

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 either* $m_1 = 1$ *or* $m_1 = 0$ *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 type of information that is decision relevant is that about x_2. Hence, if m_2 did not reveal any information about x_2, then 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 her signal. This implies that DM cannot not learn anything about the ability of the expert. Can DM learn something about the integrity of the expert when messages are uncorrelated to signals and thus do not reveal information about either x_1 or ability? 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, messages would reveal information about integrity. Accordingly, DM would retain (fire) with higher probability an expert that reports the message that is sent more often by U (B). However, this cannot be an equilibrium. Since messages do not carry information about the state, they do not affect the DM's decision about a_t. But then B would always deviate and report the message that is sent more often by U in order to increase her chances of being retained.

Lemma 5(i) implies that in all informative equilibria messages must be correlated with signals and hence with states. Since (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, this allows us to state the following Corollary to Lemma 5 (i).

Corollary 1 *In any informative equilibrium,* $\Pr(S \mid m_1 = x_1) > \alpha > \Pr(S \mid m_1 \neq x_1)$.

ii) First, consider the case in which U always reports $m_1 = 1$ regardless of s_1. Then, by Lemma 5 (i) a necessary condition for the equilibrium to be informative is that B truthfully reports at least one signal with positive probability. However, this cannot be part of the equilibrium since sending $m_1 = 0$ 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). Now consider the case in which U always reports $m_1 = 0$ regardless of s_1. Also in this case, B must truthfully report at least one signal with positive probability. It is immediate to verify that these equilibrium strategies by U and B would imply that:

$$V(0,1) > V(0,0) > V > V(1,0) = V(1,1) = 0.$$

Accordingly DM would always retain the expert after observing $m_1 = 0$, and would always replace her after observing $m_1 = 1$. Conditions (19) and (20) would therefore read:

$$(1-\delta_E) \quad \leq \quad \delta_E \pi_{2,B}(0,0) \text{ if } s_1 = 0, \tag{23}$$

$$(1-\delta_E) \quad \geq \quad \delta_E \pi_{2,B}(0,1) \text{ if } s_1 = 1, \tag{24}$$

Now notice that $V(0,1) > V(0,0)$ implies that $\pi_{2,B}(0,1) > \pi_{2,B}(0,0)$. Therefore, whenever (23) is satisfied (24) never is implying that B always sends message 0. Likewise, when (24) is satisfied (23) never is implying that B always sends message 1. In neither of these cases is Lemma 5 (i) satisfied implying that these equilibria cannot be informative.

iii) We first show that it cannot be true that both (21) and (22) are satisfied with equality, implying that it cannot be true that both $\lambda_{U,1} \in (0,1)$ and $\lambda_{U,0} \in (0,1)$. Note that both (21) and (22) 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 (21) and (22) 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 $V(m_1, x_1) > V$ so that it would be optimal for DM to retain the expert. To see this, notice that if the equilibrium is informative and $\iota(m_1, x_1) = 0$ for all (m_1, x_1), B has no career concerns and thus always sends $m_1 = 1$, but then it must be that U truthfully reports her signals with positive probability thereby sending $m_1 = 0$ with positive probability. This implies that $V(0,0) > V$.

We now prove that in all the other cases, (21) and (22) are never jointly satisfied. 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 $\Pr(S \mid m_1, x_1) = \Pr(S \mid U, m_1, x_1)\Pr(U \mid m_1, x_1) + \Pr(S \mid B, m_1, x_1)\Pr(B \mid m_1, x_1)$ and $V(m_1, x_1) = \Pr(S \mid U, m_1, x_1)\Pr(U \mid m_1, x_1)$. Since by Corollary 1 $\Pr(S \mid 1,1) > \Pr(S \mid 0,1)$, the fact that $V(1,1) = V(0,1)$ implies that $\Pr(S \mid B,1,1)\Pr(B \mid 1,1) > \Pr(S \mid B,0,1)\Pr(B \mid 0,1)$ which is equivalent to $\Pr(B \mid S,1,1)\Pr(S \mid 1,1) > \Pr(B \mid S,0,1)\Pr(S \mid 0,1)$. However, this in turn implies that $\Pr(U \mid S,1,1)\Pr(S \mid 1,1) < \Pr(U \mid S,0,1)\Pr(S \mid 0,1)$ which is equivalent to saying that $V(1,1) < V(0,1)$ which is a contradiction.

The only other possible case is when $V(1,1)\iota(1,1) = V(0,1)\iota(0,1) = 0$. Since in an informative equilibrium $V(m_1, x_1) > 0$ for all (m_1, x_1), this case occurs if and only if $\iota(1,1) = \iota(0,1) = 0$. Now, note that if $\iota(1,1) = 0$, it must be that $V(1,1) < V$. Since we can write $V(1,1) = \Pr(U \mid S,1,1)\Pr(S \mid 1,1)$, and since by Corollary 1 $\Pr(S \mid 1,1) > \alpha$, in order for $V(1,1) < V = \gamma\alpha$ it must be that $\Pr(U \mid S,1,1) < \gamma$, therefore implying that $m_1 = 1$ is a negative signal of integrity. This naturally implies that also $\iota(1,0) = 0$. But then, if also $\iota(0,0) = 0$, 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 (21) is always strictly positive.

The same line of reasoning applies to show that it cannot be that both (19) and (20) 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) To prove this, we 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 (20) 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 (22) is satisfied with equality, which further implies that the expression above is always strictly positive whenever (22) 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 $m_1 = 0$ is a positive signal for integrity and hence $m_1 = 1$ 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 (22).

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 (22).

This implies that (20) 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:

$$\widehat{\gamma}(m_1, x_1) \equiv \Pr(U \mid m_1, x_1) = \frac{\gamma \Pr(m_1 \mid U, x_1)}{\gamma \Pr(m_1 \mid U, x_1) + (1-\gamma)\Pr(m_1 \mid B, x_1)}. \tag{25}$$

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, we know by lemma 5(ii) that U must truthfully report each signal with positive probability. This means that U's messages are correlated with U's signals and hence with state x_1, implying that $\Pr(S \mid U, m_1 = x_1) > \Pr(S \mid U, m_1 \neq x_1)$. At the same time, since both U and B use the same strategy, we have that $\Pr(m_1 \mid U, x_1) = \Pr(m_1 \mid B, x_1)$ 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 (21) and (22) would always be satisfied with strict inequality, implying that U would always truthfully reveal all her signals (contradicting our initial assumption that U misreports signal $1-s'$). Having shown that U and B must misreport $1-s'$ with different probabilities, let us now assume without loss of generality that U reports $1-s'$ with higher probability than B. That is, let us assume that $\lambda_{U,s'} = \lambda_{B,s'} = 1$ 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 preference type $i = U, B$, both B and U have the same probability $\frac{1}{2}$ of observing a given signal. It follows that U reports message $m_1 = 1-s'$ more frequently than B, and message $m_1 = s'$ less frequently than B. Formally, $\Pr(m_1 = 1-s' \mid U, x_1) > \Pr(m_1 = 1-s' \mid B, x_1)$ and $\Pr(m_1 = s' \mid U, x_1) < \Pr(m_1 = s' \mid B, x_1)$, which implies $\widehat{\gamma}(m_1 = 1-s', x_1) > \gamma > \widehat{\gamma}(m_1 = s', x_1)$. 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, assume that $\lambda_{U,s'} = \lambda_{U,1-s'} = 1$ 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 prove part (ii) of Proposition 1. Consider the updates on the probability of being smart conditional on the expert being unbiased and on $m_1 = x_1$: updates on ability when the unbiased expert's message turns out to be correct:

$$\begin{aligned}\widehat{\alpha}(U,0,0) &\equiv \Pr(S \mid U,0,0) = \frac{\alpha\left[p\lambda_{U,0}+(1-p)(1-\lambda_{U,1})\right]}{\lambda_{U,0}q+(1-q)(1-\lambda_{U,1})},\\ \widehat{\alpha}(U,1,1) &\equiv \Pr(S \mid U,1,1) = \frac{\alpha\left[p\lambda_{U,1}+(1-p)(1-\lambda_{U,0})\right]}{\lambda_{U,1}q+(1-q)(1-\lambda_{U,0})},\end{aligned}$$

In an equilibrium in which U truthtells, $\lambda_{U,0} = \lambda_{U,1} = 1$ and hence $\widehat{\alpha}(U,0,0) = \widehat{\alpha}(U,1,1) = \frac{\alpha p}{q}$. In an equilibrium in which U lies, $\lambda_{U,0} \leq 1$ and $\lambda_{U,1} \leq 1$ with at least one strict inequality. Since both $\widehat{\alpha}(U,0,0)$ and $\widehat{\alpha}(U,1,1)$ are strictly increasing in $\lambda_{U,0}$ and $\lambda_{U,1}$, for any equilibrium in which U lies it is then true that $\widehat{\alpha}(U,0,0) \leq \frac{\alpha p}{q}$ and $\widehat{\alpha}(U,1,1) \leq \frac{\alpha p}{q}$ with at least one strict inequality. A similar logic applies to the case in which the unbiased expert sends a message that turns out to be incorrect, i.e., $\widehat{\alpha}(U,1,0)$ and $\widehat{\alpha}(U,0,1)$ which are strictly decreasing in $\lambda_{U,0}$ and $\lambda_{U,1}$.

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)

In a TT equilibrium, the value function representing the value of an expert in period 2 reads:

$$V^{TT}(m_1,x_1) = \begin{cases} \widehat{\gamma}^{TT}(0,0)\widehat{\alpha}^{TT}(U,0,0) & \text{if } m_1 = x_1 = 0,\\ \widehat{\gamma}^{TT}(1,1)\widehat{\alpha}^{TT}(U,1,1) & \text{if } m_1 = x_1 = 1,\\ \widehat{\gamma}^{TT}(0,1)\widehat{\alpha}^{TT}(U,0,1) & \text{if } m_1 = 0 \neq x_1 = 1,\\ \widehat{\gamma}^{TT}(1,0)\widehat{\alpha}^{TT}(U,1,0) & \text{if } m_1 = 1 \neq x_1 = 0.\end{cases}$$

It is straightforward to verify that:

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

where $\underline{\alpha} = \frac{(1-p)\alpha}{1-q}$ and $\overline{\alpha} = \frac{p\alpha}{q}$. This in turn implies that:

$$\underline{V} \equiv V^{TT}(0,1) = V^{TT}(1,0) < V < V^{TT}(1,1) = V^{TT}(0,0) \equiv \overline{V}. \tag{26}$$

where $\underline{V} = \frac{(1-p)\alpha\gamma}{1-q}$ and $\overline{V} = \frac{p\alpha\gamma}{q}$.

DM**'s strategy**. From (26), 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:

$$\imath(m_1, x_1) = \begin{cases} 1 \text{ if } m_1 = x_1, \\ 0 \text{ if } m_1 \neq x_1. \end{cases} \tag{27}$$

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 (19) gives the condition for B to truthfully report $m_1 = 0$ after observing $s_1 = 0$. By using the expression of B's continuation values, we can write (19) as follows:

$$\begin{aligned} &(1-\delta_E)\, a(0) + \delta_E \sum\nolimits_{x_1} \Pr(x_1 \mid s_1 = 0)\, [V(0,x_1) + 1]\, \imath(0,x_1) + \\ &-(1-\delta_E)\, a(1) + \delta_E \sum\nolimits_{x_1} \Pr(x_1 \mid s_1 = 0)\, [V(1,x_1) + 1]\, \imath(1,x_1) \geq 0. \end{aligned} \tag{28}$$

Now, by using (26), (27) and the fact that $\Pr(x_1 = 0 \mid s_1 = 0) = q$, condition (28) boils down to:

$$\delta_E \geq \frac{1}{(2q-1)\overline{V} + 2q} \equiv \underline{\delta}_E^{TT}. \tag{29}$$

Note that since $1/2 < q < 1$, $\underline{\delta}_E^{TT} \in (0,1)$.

U's strategy. 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 (21) gives the condition for U to truthfully report $m_1 = 0$ after observing $s_1 = 0$. By using U's continuation values, (21) can be written as:

$$\sum\nolimits_{x_1} \Pr(x_1 \mid s_1 = 0) V^{TT}(0,x_1) i(0,x_1) - \sum\nolimits_{x_1} \Pr(x_1 \mid s_1 = 0) V^{TT}(1,x_1) i(1,x_1) \geq 0. \tag{30}$$

Finally, by using (26), (27) and the fact that $\Pr(x_1 = 0 \mid s_1 = 0) = q$, condition (30) simplifies to:

$$(2q-1)\overline{V} \geq 0, \tag{31}$$

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

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)

In an MB equilibrium, the value function representing the value of an expert in period 2 reads:

$$V^{MB}(m_1, x_1) = \begin{cases} \widehat{\gamma}^{MB}(0,0)\overline{\alpha} & \text{if } m_1 = x_1 = 0, \\ \widehat{\gamma}^{MB}(1,1)\overline{\alpha} & \text{if } m_1 = x_1 = 1, \\ \widehat{\gamma}^{MB}(0,1)\underline{\alpha} & \text{if } m_1 = 0 \neq x_1 = 1, \\ \widehat{\gamma}^{MB}(1,0)\underline{\alpha} & \text{if } m_1 = 1 \neq x_1 = 0. \end{cases}$$

Note that:

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

This implies:

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

Having established this result, let us prove the existence of MB in two steps.

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

Observe that (32) implies that DM retains the expert after realization $(0,0)$ and fires her after realization $(1,0)$. This further implies that a necessary condition for the existence of our putative MB equilibrium is that the expert is retained after $(1,1)$. If not, the expert would always be fired when sending $m_1 = 1$. Hence U, whose only concern is to be retained, would never send $m_1 = 1$ (which contradicts her equilibrium strategy in MB). We now show that the condition for DM to retain the expert after $(1,1)$ is satisfied if and only if the condition for DM to fire the expert after $(0,1)$ is satisfied too. These two conditions read respectively:

$$V^{MB}(1,1) \equiv \gamma^{MB}(1,1)\overline{\alpha} > \gamma\alpha \equiv V, \tag{33}$$

$$V^{MB}(0,1) \equiv \gamma^{MB}(0,1)\underline{\alpha} < \gamma\alpha \equiv V. \tag{34}$$

By using the expressions of $\widehat{\gamma}^{MB}(1,1)$, $\overline{\alpha}$, $\widehat{\gamma}^{MB}(0,1)$ and $\underline{\alpha}$ into (33) and (34) we find that $V^{MB}(1,1) = V^{MB}(0,1) = V$ at $\lambda_{B,0} = \frac{(2p-1)(1-\alpha)}{(1+\alpha-2p\alpha)(1-\gamma)}$. This allows us to conclude that both

(33) and (34) are simultaneously satisfied if and only if $\lambda_{B,0} > \frac{(2p-1)(1-\alpha)}{(1+\alpha-2p\alpha)(1-\gamma)} \equiv \lambda'_B \in (0,1)$. Note that this further implies that the strategy of DM to retain after $m_1 = x_1$ and fire after $m_1 \neq x_1$ is the only strategy that is consistent with an MB equilibrium.

Step 2) We now show that for sufficiently high values of $\lambda_{B,0}$, U's and B's equilibrium strategies are optimal given DM's equilibrium strategy outlined in Step 1.

U's strategy. Let's first consider the case in which $s_1 = 1$. U truthfully reports signal $s_1 = 1$ if condition (22) is satisfied. Given DM's strategy, condition (22) becomes:

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

which can be written as:

$$q\widehat{\gamma}^{MB}(1,1) \geq (1-q)\widehat{\gamma}^{MB}(0,0). \tag{35}$$

where $q > 1-q$. Now note that:

i) When $\lambda_{B,0} = 0$, we have that $\widehat{\gamma}^{MB}(1,1) = 0$ and $\widehat{\gamma}^{MB}(0,0) = 1$, implying that $LHS < RHS$;

ii) When $\lambda_{B,0} = 1$, we have that $\widehat{\gamma}^{MB}(1,1) = \widehat{\gamma}^{MB}(0,0) = \gamma$, implying that $LHS > RHS$;

iii) $\widehat{\gamma}^{MB}(1,1)$ and $\widehat{\gamma}^{MB}(0,0)$ 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 $\lambda_{B,0} \in \left[\widetilde{\lambda}_B, 1\right]$, (35) is satisfied. Let's now consider the case $s_1 = 0$, for which the relevant condition for truthtelling is given by expression (21). It is immediate to note that (21) 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 (19) 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 $\Pr(x_1 = 0 \mid s_1 = 0) = q$, we can write the previous condition as:

$$\delta_E = \frac{1}{[qV^{MB}(0,0)] - (1-q)V^{MB}(1,1)] + 2q} \equiv \delta_E^{MB}(\lambda_{B,0}). \tag{36}$$

Note that since $1/2 < q < 1$ 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 (36) coincides with the RHS of (29). 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

Lemma 5 (ii) implies that U always sends both messages with positive probability in equilibrium and Lemma 5 (iii) implies that, in an informative equilibrium, U can misreport at most one signal. Hence, we can conveniently define equilibria in which U misreports in the following way:

- *Misreporting Unbiased equilibria* (MU in short): equilibria in which 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 $\lambda_{B,1} = 1$. This implies that we can further restrict our attention on the existence of the following two putative sub-classes of equilibria belonging to MU:

i) $MU(1)$: 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 $\lambda_{B,0} \in [0,1]$.

ii) $MU(0)$: 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 $\lambda_{B,0} \in [0,1]$.

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

A.5.1 $MU(1)$ 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 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 $s_1 = 0$. In an $MU(1)$ equilibrium, The value function representing the value of an expert in period 2 reads:

$$V^{MU(1)}(m_1,x_1) = \begin{cases} \widehat{\gamma}^{MU(1)}(0,0)\widehat{\alpha}^{MU(1)}(U,0,0) & \text{if } m_1 = x_1 = 0, \\ \widehat{\gamma}^{MU(1)}(1,1)\widehat{\alpha}^{MU(1)}(U,1,1) & \text{if } m_1 = x_1 = 1, \\ \widehat{\gamma}^{MU(1)}(0,1)\widehat{\alpha}^{MU(1)}(U,0,1) & \text{if } m_1 = 0 \neq x_1 = 1, \\ \widehat{\gamma}^{MU(1)}(1,0)\widehat{\alpha}^{MU(1)}(U,1,0) & \text{if } m_1 = 1 \neq x_1 = 0. \end{cases}$$

Note that:

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

This immediately implies:

$$V^{MU(1)}(1,0) < V < V^{MU(1)}(0,0). \tag{37}$$

In order to prove existence we proceed in two steps.

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

First, note that condition (37) 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 the condition for DM to retain the expert after $(1,1)$ is satisfied if and only if the condition for DM to fire the expert after $(1,0)$ is satisfied too. These two conditions read respectively:

$$V^{MU(1)}(1,1) = \widehat{\gamma}^{MU(1)}(1,1)\widehat{\alpha}^{MU(1)}(U,1,1) > \gamma\alpha \equiv V, \tag{38}$$

$$V^{MU(1)}(0,1) = \widehat{\gamma}^{MU(1)}(0,1)\widehat{\alpha}^{MU(1)}(U,0,1) < \gamma\alpha \equiv V. \tag{39}$$

By substituting the expressions of $\widehat{\gamma}^{MU(1)}(1,1)$, $\widehat{\alpha}^{MU(1)}(U,1,1)$, $\widehat{\gamma}^{MU(1)}(0,1)$, $\widehat{\alpha}^{MU(1)}(U,0,1)$ into (38) and (39), we find that $V^{MU(1)}(1,1) = V^{MU(1)}(0,1) = V$ at $\lambda_{U,1} = \frac{(1-\alpha+2p\alpha)(1-\gamma)}{2p-\gamma+\alpha\gamma-2p\alpha\gamma}$. This allows us to conclude that both (38) and (39) are simultaneously satisfied if and only if $\lambda_{U,1} > \frac{(1-\alpha+2p\alpha)(1-\gamma)}{2p-\gamma+\alpha\gamma-2p\alpha\gamma} \equiv \lambda'_{U,1}$. Note that this further implies that the strategy of DM to retain the expert after $m_1 = x_1$ and fire her after $m_1 \neq x_1$ is the only strategy that is consistent with an MU equilibrium.

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:

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

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

$$\delta_E[qV^{MU(1)}(0,0) - (1-q)V^{MU(1)}(1,1) + 2q - 1] > 1 - \delta_E. \tag{41}$$

Note that since $q > \frac{1}{2}$, if condition (40) 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 (41) is strictly increasing in δ_E. Since the RHS is always strictly decreasing in δ_E, we can conclude that if condition (40) is satisfied, then there always exists a value of δ_E above which (41) is satisfied as well.

This means that we only need to show that condition (40) 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 (40) 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 (40) 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 (40) is satisfied.

Case in which B truthfully reports $s_1 = 1$ and misreports $s_1 = 0$. This is the case in which $\lambda_{U,1} < 1$ and $\lambda_{B,0} < 1$. First note that the inequalities given by (37) continue to hold true. 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:

$$qV^{MU(1)}(1,1) = (1-q)V^{MU(1)}(0,0), \tag{42}$$

$$\delta_E q[V^{MU(1)}(0,0)+1] = (1-\delta_E) + \delta_E(1-q)[V^{MU(1)}(1,1)+1], \tag{43}$$

$$V^{MU(1)}(0,1) < V < V^{MU(1)}(1,1). \tag{44}$$

Condition **(42)** is the condition that must be satisfied for U to randomize after $s_1 = 1$. Condition **(43)** is the condition that must be satisfied in order for B to randomize after $s_1 = 0$. Finally, condition **(44)** 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 (42). Let $\lambda^*_{U,1} \in (0,1)$ $[\lambda^*_{U,1} \in (\lambda'_U, 1)]$ be the value of $\lambda_{U,1}$ that satisfies (42) 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 $V(1,1)$ and $V(0,0)$ are respectively strictly increasing and decreasing in $\lambda_{B,0}$. Hence, when $\lambda_{B,0} = 1-\varepsilon$ (with $\varepsilon > 0$) and $\lambda_{U,1} = \lambda^*_{U,1}$ we have that: $qV(1,1) < (1-q)V(0,0)$. By the proof of proposition 2 we also know that when $\lambda_{B,0} = 1-\varepsilon \geq \lambda^*_B$ and $\lambda_{U,1} = 1$ (i.e., when we are in an MB equilibrium) we have that: $qV(1,1) > (1-q)V(0,0)$. But then, when $\lambda_{B,0} = 1-\varepsilon$, by continuity there must exist a $\lambda_{U,1} \in (\lambda^*_U, 1)$ such that $qV(1,1) = (1-q)V(0,0)$.

Second, let's consider condition **(44).** 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 (43). If we solve it for δ_E we obtain:

$$\delta_E = \frac{1}{qV^{MU(1)}(0,0) - (1-q)V^{MU(1)}(1,1) + 2q}. \tag{45}$$

If (42) is satisfied, $qV^{MU(1)}(0,0) - (1-q)V^{MU(1)}(1,1)$ 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 (42) and (44) are satisfied, we can always find a value of $\delta_E \in (0,1)$ that guarantees that condition (43) 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 (45) takes in $MU(1)$.

A.5.2 $MU(0)$ Equilibria

Also for this case, we first prove that there exist $MU(0)$ equilibria where $\lambda_{B,0} = 1$ (i.e., $MU(0)$ equilibria where B truthfully reports both signals), and then prove that there also exist $MU(0)$ equilibria where $\lambda_{B,0} \in [0, 1)$ (i.e., $MU(0)$ equilibria where B truthfully reports $s_1 = 1$ and misreports $s_1 = 0$).

Case in which B truthfully reports both $s_1 = 1$ and $s_1 = 0$. The value function representing the value of an expert in period 2 in an $MU(0)$ equilibrium reads:

$$V^{MU(0)}(m_1, x_1) = \begin{cases} \widehat{\gamma}^{MU(0)}(0,0)\widehat{\alpha}^{MU(0)}(U,0,0) & \text{if } m_1 = x_1 = 0, \\ \widehat{\gamma}^{MU(0)}(1,1)\widehat{\alpha}^{MU(0)}(U,1,1) & \text{if } m_1 = x_1 = 1, \\ \widehat{\gamma}^{MU(0)}(0,1)\widehat{\alpha}^{MU(0)}(U,0,1) & \text{if } m_1 = 0 \neq x_1 = 1, \\ \widehat{\gamma}^{MU(0)}(1,0)\widehat{\alpha}^{MU(0)}(U,1,0) & \text{if } m_1 = 1 \neq x_1 = 0. \end{cases}$$

Note that:

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

This immediately implies:

$$V^{MU(0)}(1,1) > V > V^{MU(0)}(0,1).$$

This means that 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 $m_1 = 0$ (which contradicts her equilibrium strategy). Therefore, existence requires that:

$$V^{MU(0)}(0,0) > V. \tag{46}$$

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 (46) is satisfied if and only if $\lambda_{U,0} > \frac{(1-\alpha+2p\alpha)(1-\gamma)}{2p-\gamma+\alpha\gamma-2p\alpha\gamma} \equiv \lambda'_{U,0}$; ii) 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 $\lambda_{U,0} > \lambda'_{U,0}$ provided that α is sufficiently small. Hence, also $MU(0)$ 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 (19) must be satisfied with strict inequality. Since the equilibrium behavior of U implies that $qV^{MU(0)}(0,0) = (1-q)V^{MU(0)}(1,1)$, (19) 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 $s_1 = 0$, (21) 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 $\widehat{\alpha}^{MU(0)}(U,0,0) > \widehat{\alpha}^{MU(0)}(U,1,1)$. Hence, to have that $V^{MU(0)}(1,1) > V^{MU(0)}(0,0)$, it must be that $\widehat{\gamma}(1,1)^{MU(0)} > \widehat{\gamma}^{MU(0)}(0,0)$. 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 $\lambda_{U,0} < \lambda_{B,0}$ (i.e. the unbiased expert must lie more than the biased one).

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 Proof of Proposition 4

Since DM replaces the incumbent with a new expert after a mistake, we have that:

$$\begin{aligned}E_0^{MB}[R_2] &= \Pr(0,0|MB)\gamma^{MB}(0,0)\alpha^{MB}(0,0) + \Pr(1,1|MB)\gamma^{MB}(1,1)\alpha^{MB}(1,1) + \\ &\quad + [\Pr(0,1|MB) + \Pr(1,0|MB)]\,\gamma\alpha = \\ &= \Pr(0,0|MB)\gamma^{MB}(0,0)\overline{\alpha} + \Pr(1,1|MB)\gamma^{MB}(1,1)\overline{\alpha} + \\ &\quad \Pr(0,1|MB)\gamma\alpha + \Pr(1,0|MB)\gamma\alpha = \\ &= \tfrac{1}{2}q\gamma\overline{\alpha} + \tfrac{1}{2}q\gamma\overline{\alpha} + [\Pr(0,1|MB) + \Pr(1,0|MB)]\,\gamma\alpha = \\ &= q\gamma\overline{\alpha} + [\Pr(0,1|MB) + \Pr(1,0|MB)]\,\gamma\alpha.\end{aligned}$$

Now let's compare this expression with the following one arising in a TT equilibrium:

$$\begin{aligned}E_0^{TT}[R_2] &= \Pr(0,0|TT)\gamma^{TT}(0,0)\alpha^{TT}(U,0,0) + \Pr(1,1|TT)\gamma^{TT}(1,1)\alpha^{TT}(U,1,1) + \\ &\quad + [\Pr(0,1|TT) + \Pr(1,0|TT)]\,\gamma\alpha = \\ &= \tfrac{1}{2}q\gamma\overline{\alpha} + \tfrac{1}{2}q\gamma\overline{\alpha} + [\Pr(0,1,TT) + \Pr(1,0,TT)]\,\gamma\alpha = \\ &= q\gamma\overline{\alpha} + [\Pr(0,1,TT) + \Pr(1,0,TT)]\,\gamma\alpha.\end{aligned}$$

Note that:

$$\begin{aligned}\Pr(0,1|MB) + \Pr(1,0|MB) &= \frac{1}{2}[1-\alpha(2p-1)(1-(1-\gamma(1-\lambda_{B,0})] > \\ \frac{1}{2}[1-\alpha(2p-1)] &= \Pr(0,1|TT) + \Pr(1,0|TT),\end{aligned}$$

for any $\lambda_{B,0} \in [0,1)$. It follows that $E_0^{MB}[R_2] > E_0^{TT}[R_2]$.

A.7 Proof of Proposition 5

In order to determine when MU can improve sorting with respect to TT we first consider the case of $MU(1)$ equilibria and then the case of $MU(0)$ equilibria.

$MU(1)$ **Equilibria (**U **Lies after** $s_1 = 1$ **and** B **Lies after** $s_1 = 0$**)** Since DM replaces the incumbent with a new expert after a mistake, we have that:

$$\begin{aligned}E_0^{MU(1)}[R_2] &= \Pr(0,0|MU(1))\widehat{\gamma}^{MU(1)}(0,0)\widehat{\alpha}^{MU(1)}(U,0,0) \\ &\quad + \Pr(1,1|MU(1))\widehat{\gamma}^{MU(1)}(1,1)\widehat{\alpha}^{MU(1)}(U,1,1) + \\ &\quad + [\Pr(0,1,MU(1))\Pr(1,0,MU(1))]\,\gamma\alpha.\end{aligned}$$

By using the equilibrium values of $\Pr(m_1, x_1|MU(1))$, $\gamma^{MU(1)}(m_1, x_1)$ and $\alpha^{MU(1)}(U, m_1, x_1)$, $E_0^{MU(1)}[R_2]$ can be written as follows:

$$E_0^{MU(1)}[R_2] = \frac{1}{2}\gamma\alpha\left\{2 + (2p-1)\lambda_{U,1} - (2p-1)\alpha\left[(1-\gamma)\lambda_{B,0} + \gamma\lambda_{U,1})\right]\right\}.$$

Now let's compare this last expression with the following one arising in a TT equilibrium:

$$E_0^{TT}[R_2] = \Pr(0,0|TT)\widehat{\gamma}^{TT}(0,0)\widehat{\alpha}^{TT}(U,0,0) + \Pr(1,1|TT)\widehat{\gamma}^{TT}(1,1)\widehat{\alpha}^{TT}(U,1,1) + \\ + \left[\Pr(0,1|TT) + \Pr(1,0|TT)\right]\gamma\alpha.$$

By using the equilibrium values of $\Pr(m_1, x_1|TT)$, $\gamma^{TT}(m_1, x_1)$ and $\alpha^{TT}(U, m_1, x_1)$, $E_0^{TT}[R_2]$ can be written as follows:

$$E_0^{TT}[R_2] = \frac{1}{2}\gamma\alpha\left[2 + (2p-1)(1-\alpha)\right].$$

It is easy to verify that $E_0^{MU(1)}[R_2] \geq E_0^{TT}[R_2]$ if $\lambda_{U,1} \geq \frac{1-\alpha+\alpha\lambda_{B,0}-\alpha\gamma\lambda_{B,0}}{1-\alpha\gamma}$. Now note that $\frac{1-\alpha+\alpha\lambda_{B,0}-\alpha\gamma\lambda_{B,0}}{1-\alpha\gamma} > \lambda_{B,0}$. This implies that $E_0^{MU(1)}[R_2] \geq E_0^{TT}[R_2]$ only if $\lambda_{U,1} > \lambda_{B,0}$. Hence, a necessary condition for $MU(1)$ to dominate TT is that the B lies more than U.

We now prove that $MU(1)$ equilibria with such features exist. In particular, we prove that there exist $MU(1)$ equilibria characterized by $\lambda_{U,1} \to 1$ and $\lambda_{B,0} < 1$. Since it is true that $\frac{1-\alpha+\alpha\lambda_{B,0}-\alpha\gamma\lambda_{B,0}}{1-\alpha\gamma} < 1$ for any $\lambda_{B,0} < 1$, these $MU(1)$ equilibria improve sorting over TT. To do this we proceed as follows:

Step 1: We can iterate the procedure used in the proof of Proposition 3 to show the existence of $MU(1)$ in which B misreports when receiving $s_1 = 0$. Thus starting from the $MU(1)$ equilibrium in which $\lambda_{B,0} = 1 - \varepsilon$ and $\lambda_{U,1} = \lambda_{U,1}^{**} \in (\lambda_{U,1}^{*}, 1)$, if we further reduce $\lambda_{B,0}$ by ε so that $\lambda_{B,0} = 1 - 2\varepsilon$ when $\lambda_{U,1} = \lambda_{U,1}^{**}$, we have that $qV(1,1) < (1-q)V(0,0)$. We know that if $\lambda_{B,0} = 1 - 2\varepsilon > \lambda_B^{*}$ (with λ_B^{*} being defined in the proof of proposition 2), there exists an MB equilibrium characterized by $\lambda_{B,0} = 1 - 2\varepsilon$ and $\lambda_{U,1} = 1$, in which case $qV(1,1) > (1-q)V(0,0)$. Therefore by continuity there always exists a $\lambda_{U,1} \in (\lambda_{U,1}^{**}, 1)$ for which a new $MU(1)$ equilibrium exists for $\lambda_{B,0} = 1 - 2\varepsilon$.

Step 2: We can iterate the procedure outlined in Step 1 until we reach $\lambda_{B,0} = \lambda_B^{*} + \varepsilon$ where we can define $\widehat{\lambda}_{U,1} < 1$ as the corresponding probability that satisfies $MU(1)$ when $\lambda_{B,0} = \lambda_B^{*} + \varepsilon$. Here, based on the proof of Proposition 2, we have two possible cases:

a) $\lambda_{B,0}^{*} = \lambda_B' > \widetilde{\lambda}_B$ in which the binding condition for the existence of MB is determined by the hiring strategy of the DM;

b) $\lambda_B^{*} = \widetilde{\lambda}_B > \lambda_B'$ in which the binding condition for the existence of MB is determined by U's truthtelling condition.

Before considering each case, we prove the following Lemma which will also be useful for proving Proposition 6:

Lemma 6 *$MU(1)$ equilibria with higher levels of $\lambda_{U,1}$ are characterized by lower levels of $\lambda_{B,0}$ and lower levels of δ_E.*

Proof

To prove this, first note that by Steps 1 and 2 it follows that $MU(1)$ equilibria characterized by higher values of $\lambda_{U,1}$ have lower values of $\lambda_{B,0}$. Recall from the proof of Proposition 3, that there exists an $MU(1)$ equilibrium in which $\lambda_{B,0} = 1$ and $\lambda_{U,1} = \lambda^*_{U,1}$. Denote the corresponding value of $V(1,1)$ in this equilibrium with $V^*(1,1)$. Now, denote the equilibrium values of $V(1,1)$ when $\lambda_{B,0} = 1-\varepsilon$ and $\lambda_{U,1} = \lambda^*_{U,1}$ with $V^-(1,1)$, and when $\lambda_{B,0} = 1-\varepsilon$ and $\lambda_{U,1} = 1$ with $V^+(1,1)$. Since $\lambda^*_{U,1} < 1$,if we choose a positive value of ε sufficiently close to 0, $V^-(1,1) \rightarrow V(1,1)$ and $V^+(1,1) > V(1,1)$. This implies that for $\lambda^{**}_{U,1} \in (\lambda^*_{U,1}, 1)$ that satisfies the MU equilibrium for $\lambda_{B,0} = 1-\varepsilon$, the corresponding value $V^{**}(1,1)$ is such that $V^{**}(1,1) > V^*(1,1)$. Iterating this procedure for all $MU(1)$ equilibria as in Step 2, this proves that $MU(1)$ with higher values of $\lambda_{U,1}$, are characterized by lower values $\lambda_{B,0}$, and higher values of $V(1,1)$.

Now, let us denote with $\lambda^{eq}_{U,1}$ the equilibrium value of$lambda_{U,1}$, and with $V^{eq}(m_1,x_1)$ the equilibrium value of $V(m_1,x_1)$ in an $MU(1)$ equilibrium. Considering any $MU(1)$ equilibrium with $\lambda^{eq}_{U,1} < 1$, if we take the total differential of the condition that leads U to randomize with respect to $\lambda^{eq}_{U,1}$ (i..e. moving from one $MU(1)$ equilibria to the other) it must be that:

$$q\frac{\partial V^{eq}(1,1)}{\partial \lambda^{eq}_{U,1}} - (1-q)\frac{\partial V^{eq}(0,0)}{\partial \lambda^{eq}_{U,1}} = 0 \rightarrow \frac{q}{(1-q)} = \frac{\frac{\partial V^{eq}(0,0)}{\partial \lambda^{eq}_{U,1}}}{\frac{\partial V^{eq}(1,1)}{\partial \lambda^{eq}_{U,1}}}. \tag{47}$$

In other words, since $\frac{q}{(1-q)} > 1$, the equilibrium variation in $V^{eq}(0,0)$ must be of the same sign and of a greater magnitude with respect to the variation in $V^{eq}(1,1)$ as we move from an $MU(1)$ equilibrium with a lower $\lambda^{eq}_{U,1}$(and higher $\lambda^{eq}_{B,0}$) to one with higher $\lambda^{eq}_{U,1}$(and lower $\lambda^{eq}_{B,0}$).

We now show that in $MU(1)$ equilibria, higher levels of $\lambda_{U,1}$ are associated to lower levels of δ_E, i.e. $\frac{\partial \delta_E}{\partial \lambda^{eq}_{U,1}} < 0$. Consider (45). It is straightforward to see that:

$$sign\left(\frac{\partial \delta_E}{\partial \lambda^{eq}_{U,1}}\right) = -sign\left(\frac{\partial[qV^{eq}(0,0) - (1-q)V^{eq}(1,1)]}{\partial \lambda^{eq}_{U,1}}\right).$$

To show that $\frac{\partial \delta_E}{\partial \lambda^{eq}_{U,1}} < 0$, it is therefore sufficient to prove that $\frac{\partial[qV^{eq}(0,0)-(1-q)V^{eq}(1,1)]}{\partial \lambda^{eq}_{U,1}} > 0$ or, equivalently, that:

$$q/(1-q) > \frac{\frac{\partial V^{eq}(1,1)}{\partial \lambda_{U,1}^{eq}}}{\frac{\partial V^{eq}(0,0)}{\partial \lambda_{U,1}^{eq}}}.$$

Since $q > (1-q)$ and $\frac{\partial V^{eq}(0,0)}{\partial \lambda_{U,1}^{eq}} > \frac{\partial V^{eq}(1,1)}{\partial \lambda_{U,1}^{eq}}$, the above inequality is always satisfied.

Let us now go back to our cases a) and b).

Case a)

Suppose $\lambda_{B,0} = \lambda'_B + \varepsilon$. This implies that there exists an MB equilibrium such that $V(0,1) < \alpha\gamma < V(1,1)$, and there also exists an $MU(1)$ with $\lambda_{U,1} = \widehat{\lambda}_{U,1}$. Now, if we further reduce $\lambda_{B,0}$ from $\lambda'_B + \varepsilon$ to λ'_B, we have that $V(1,1) = \alpha\gamma$ and $V(0,1) = \alpha\gamma$. Since $\lambda'_B > \widetilde{\lambda}_B$, we know that MB continues to hold. In order to obtain an $MU(1)$ equilibrium starting from an MB with values of $\lambda_{B,0} \leq \lambda'_B$, we need to reduce $\lambda_{U,1}$ below 1. However, if we reduce $\lambda_{U,1}$ below 1 when $\lambda_{B,0} = \lambda'_B$, $V(1,1)$ decreases and $V(0,1)$ increases (since $V(1,1)$ and $V(0,1)$ are respectively increasing and decreasing in $\lambda_{U,1}$) implying that $V(1,1) < \alpha\gamma$ and $V(0,1) > \alpha\gamma$, and the equilibrium condition for the DM is violated. Thus $MU(1)$ equilibria may exist only for $\lambda_{B,0} > \lambda'_B$ and $\lambda_{U,1} = \widehat{\lambda}_{U,1} < 1$, which does not guarantee the existence of MU^* equilibria in this case.

Case b)

Iterating the procedure outlined in Step 2, as $\lambda_{B,0}$ tends to $\widetilde{\lambda}_B$ from above we have that by the definition of $\widetilde{\lambda}_B$, $qV(1,1) = (1-q)V(0,0)$ for $\lambda_{U,1} = 1$ and $qV(1,1) < (1-q)V(0,0)$ for $\lambda_{U,1} < 1$. By continuity this implies that there always exists an $MU(1)$ with $\lambda_{U,1} \to 1$ and $\lambda_{B,0}$ tending to $\widetilde{\lambda}_B$ from above. This guarantees that MU^* equilibria exist in this case, and by Lemma 6 this also implies that for $\lambda_{B,0} < \widetilde{\lambda}_B$ no $MU(1)$ exist.

The following corollary is an immediate result of the proof above:

Corollary 2 *$MU(1)$ equilibria exist if and only if $\lambda_{B,0} \in [\lambda_B^{MU}, 1]$ where $\lambda_B^{MU} \in (\lambda'_B, 1)$*

This corollary implies that $MU(1)$ equilibria exist for a subset of values of $\lambda_{B,0}$ for which MB equilibria exist with the lower bound of $\lambda_{B,0}$ being strictly greater than λ'_B.

$MU(0)$ Equilibria (Both U and B Lie after $s_1 = 0$). Since DM replaces the incumbent with a new expert after a mistake, we have that:

$$\begin{aligned} E_0^{MU(0)}[R_2] = & \Pr(0,0|MU(0))\widehat{\gamma}^{MU(0)}(0,0)\widehat{\alpha}^{MU(0)}(U,0,0) \\ & + \Pr(1,1|MU(0))\widehat{\gamma}^{MU(0)}(1,1)\widehat{\alpha}^{MU(0)}(U,1,1)+ \\ & + [\Pr(0,1,MU(0))\Pr(1,0,MU(0))]\,\gamma\alpha \end{aligned}$$

By using the equilibrium values of $\Pr(m_1, x_1|MU(0))$, $\widehat{\gamma}^{MU(0)}(m_1, x_1)$ and $\widehat{\alpha}^{MU(0)}(U, m_1, x_1)$, $E_0^{MU(0)}[R_2]$ can be written as follows:

$$E_0^{MU(0)}[R_2] = \frac{1}{2}\alpha\gamma\left\{2 + (2p-1)\lambda_{U,0} + (2p-1)\alpha\left[(1-\gamma)\lambda_{B,0} + \gamma\lambda_{U,0})\right]\right\}$$

Now note that the expression of $E_0^{MU(0)}[R_2]$ has the same form of that of $E_0^{MU(1)}[R_2]$. Hence, as it can be easily verified, it is true that $E_0^{MU(0)}[R_2] \geq E_0^{TT}[R_2]$ if and only if $\lambda_{U,0} \geq \frac{1-\alpha+\alpha\lambda_{B,0}-\alpha\gamma\lambda_{B,0}}{1-\alpha\gamma} > \lambda_{B,0}$. So, we have again that $E_0^{MU(0)} \geq E_0^{TT}$ only if $\lambda_{U,0} > \lambda_{B,0}$. Put differently, a necessary condition for $MU(0)$ to dominate TT is that the biased lies more than the unbiased. In this respect, we note that, as shown in the proof of proposition 4, $MU(0)$ equilibria exist only if $\lambda_{U,0} < \lambda_{B,0}$, implying that it is never the case that they can improve sorting with respect to TT.

A.8 Proof of Proposition 6

Consider the case in which $\delta_E = \underline{\delta}_E^{TT}$. Then, a TT equilibrium exists in which $\lambda_{B,0} = 1$, $\lambda_{U,1} = 1$, and the following conditions are satisfied:

i) $V(1,1) > V > V(0,1)$

ii) $qV(1,1) > (1-q)V(0,0)$

iii) $\underline{\delta}_E^{TT} q\left[V(0,0)+1\right] = 1 - \underline{\delta}_E^{TT} + \underline{\delta}_E^{TT}(1-q)\left[V(1,1)+1\right]$

Now suppose we want to construct an MU equilibrium in which B tells the truth, which we know exists by Proposition 3. Accordingly, let us decrease $\lambda_{U,1}$ below 1. To have the MU equilibrium described above, it must be that condition (ii) is satisfied with equality, while the LHS of (iii) must be weakly larger than the RHS. When $\lambda_{U,1}$ decreases, $V(1,1)$ decreases too. Now, consider condition (ii). There are two cases in which it can be satisfied with equality: 1) $V(0,0)$ increases; 2) $V(0,0)$ decreases less $V(1,1)$. Note that in both cases, the LHS of (iii) becomes larger than the RHS. This implies that for $\delta_E = \underline{\delta}_E^{TT}$, the LHS is strictly greater than the RHS. Hence, the value of δ_E that supports an MU equilibrium in which B truthtells is strictly lower than $\underline{\delta}_E^{TT}$.

By Lemma 6 we know that $MU(1)$ equilibria characterized by lower $\lambda_{B,0}$ are necessarily characterized by lower values of δ_E. Also by Corollary 2, we know that $MU(1)$ equilibria exist for $\lambda_{B,0} \geq \lambda_B^*$, and finally by Proposition 2 we know that $\underline{\delta}_E^{MB} < \overline{\delta}_E^{MB} = \underline{\delta}_E^{TT}$. Therefore, it follows that all the $MU(1)$ equilibria in which B misreports exist for $\underline{\delta}_E^{MB} < \delta_E \in (\underline{\delta}_E^{MB}, \underline{\delta}_E^{TT})$. This proves that MB and MU^* may exist for intermediate values of δ_E that are strictly lower than the values of δ_E for which TT exists.

A.9 Non Existence of Truthtelling when the Horizon is Infinite

Consider the same stage-game described in Section 2 with the following changes. First, there is an infinite number of periods and we denote each period with t. Second, in each period in which an expert provides a recommendation, she receives a fixed fee w. Finally, in each period in which the the decision maker (DM) chooses $a_t = 1$, the biased expert receives an extra payoff $0 < y < \infty$.

We now proof that there cannot exist Markovian equilibria in which both biased (B) and unbiased experts (U) truthtell. We denote such equilibria with TT. Let ι_t represents the choice that DM makes at the end of period t to keep ($\iota_t = 1$) the incumbent expert for advice in period $t+1$ or fire ($\iota_t = 0$) and replace her with a new expert (a "challenger"). The utility functions of U and B read as follows:

$$U_U = \Sigma_{t=0}^{\infty} \delta_E^t \iota_t w,$$

$$U_B = \Sigma_{t=0}^{\infty} \delta_E^t \iota_t (y a_t + w).$$

The utility function of DM is given by:

$$U_{DM} = \Sigma_{t=0}^{\infty} \delta_{DM}^t R(a_t, x_t).$$

A Markov expert strategy is a pair (σ_U, σ_B) where $\sigma_I(s, \underline{r}_E, \underline{r}_{DM})$ is the probability of sending message 1, if the expert is of type $I \in \{U, B\}$, receives signal s and enters a given period with reputation vectors $\underline{r}_E$ and $\underline{r}_{DM}$. These vectors respectively represent the DM's probability distribution in relation to the expert's type ($\underline{r}_{DM}$), which is common knowledge, as well as the expert's probability distribution on her own type ($\underline{r}_E$). Since there are two dimensions of uncertainty, the support for $\underline{r}_{DM}$ is given by the four possible types (i.e., US, UD, BS, BD). Notice also, that the although the expert knows her integrity, what she learns about her ability may differ from what DM learns whenever there is some misreporting on the equilibrium path. A DM strategy is a pair (a, ι). $a(m, \underline{r}_{DM})$ is the decision maker's action if he receives message m, and has beliefs on the incumbent expert's type given by vector $\underline{r}_{DM}$, while $\iota(\underline{r}_{DM}, m, x)$ is his choice of keeping the incumbent expert with reputation vector $\underline{r}_{DM}$, that sends message m when the state of the world is x.

A Markov equilibrium is characterized by a strategy profile $(\sigma_U; \sigma_B; (a, \iota))$ and value functions V_U, V_B and V_{DM} such that 1) DM strategy (a, ι) is optimal given $(\sigma_U; \sigma_B)$ and V_{DM} after every history; 2) $(\sigma_U; \sigma_B)$ maximizes current plus reputational payoff (given by V_U, V_B) after every history; 3) value functions are generated by strategy profile $(\sigma_U; \sigma_B; (a, \iota))$.

In a TT equilibrium, B and U behave in the same way. Hence, uncertainty is only about the expert being smart or dumb. Accordingly, only the reputation for being smart matters, and therefore the reputation vectors simplify to scalars, and it follows that $r_{DM} = r_E \equiv \alpha$.[25]

We assume that the game starts in period $t = 1$ and we let α_0 denote the prior reputation of the incumbent expert that is hired at the beginning of $t = 1$. We are interested in the incentives of the incumbent to truthtell in each period in which he is retained and thus called upon to provide advice.

The DM starts each period t with belief α_{t-1} about the incumbent being smart. At the end of period t, after observing the recommendation of the incumbent (m_t) and state of the world (x_t), he updates his beliefs α_t about the incumbent being smart.

Hence, in a TT equilibrium we have that:

$$\alpha_t \equiv \alpha_t(\alpha_{t-1}, m_t, x_t) = \begin{cases} \frac{\alpha_{t-1}p}{\alpha_{t-1}p+(1-\alpha_{t-1})\frac{1}{2}} \equiv \alpha_t(\alpha_{t-1})^+ > \alpha_{t-1} \text{ if } m_t = x_t \\ \frac{\alpha_{t-1}(1-p)}{\alpha_{t-1}(1-p)+(1-\alpha_{t-1})\frac{1}{2}} \equiv \alpha_t(\alpha_{t-1})^- < \alpha_{t-1} \text{ if } m_t \neq x_t \end{cases}$$

Let's denote with $V^{DM}(\alpha_t)$ the value function of the DM at the end of period t with an incumbent of reputation α_t. In a TT equilibrium, the higher α_t, the more valuable the information of the expert. Hence, $V^{DM}(\alpha_t)$ is strictly increasing in α_t.

At the end of each period t, DM decides whether to retain the incumbent or to hire a challenger for advice in period $t+1$. We assume that the initial reputation of the challenger is equal to α_0. Hence, DM retains the incumbent if and only if $V^{DM}(\alpha_t) > V^{DM}(\alpha_0)$. Since $V^{DM}(\cdot)$ is strictly increasing, this is equivalent to $\alpha_t > \alpha_0$. Note that for any α_0 and $p \in \left(\frac{1}{2}, 1\right)$, there always exists a threshold $\widetilde{\alpha} \in (\alpha_0, 1)$ such that if $\alpha_{t-1} > \widetilde{\alpha}$, we have that $\alpha_t(\alpha_{t-1})^- > \alpha_0$.[26]That is, the expert is retained to provide advice in $t + 1$ even if she made a mistake in t.

Let $V^B(\alpha_t)$ denote the value function of a biased expert with reputation α_t at the end of a given period t. It is straightforward to observe that this value function is increasing in α_t.[27] Now, consider a history of our putative TT equilibrium that led to period t with

[25]The first equivalence follows from the fact that in a TT equilibrium there is no difference between what an expert and DM learn about the ability of the expert, and the second one follows from the observation that integrity does not play a role since there is no difference between the strategies of U versus B.

[26]It is straightforward to show show that $\widetilde{\alpha} \equiv \frac{\alpha_0}{2-\alpha_0-2p+2\alpha_0 p}$.

[27]Note that the value function of an expert $V^i(\alpha_t)$ with $i = B, U$ is strictly increasing in his reputation α_t. To see this, consider an expert that has developed a reputation α_{t-1} that is just above this threshold. What are the incentives of this expert in t considering that she will be retained anyway? Making a mistake will reduce the reputation with which she will end t to a level such that, a mistake in $t + 1$ will imply dismissal. On the other hand, making a correct recommendation in t will further increase the reputation to a level such that she will be retained in $t + 1$ even after a mistake. Hence, the higher the reputation, the higher the probability that the expert will be retained for a given number of periods, and therefore the higher the continuation value of the expert.

a sufficient number of correct recommendations such that $\alpha_t \to 1$. In this case, B will be kept at the end of period $t+1$ independently from whether her recommendation turns out to be correct. We then show that truthtelling cannot be an equilibrium. The condition for B to TT in period $t+1$ when receiving $s_{t+1}=0$ can be written in the following way:

$$\delta_E[pV^B(\alpha_{t+1}(\alpha_t)^+) - (1-p)V^B(\alpha_{t+1}(\alpha_t)^-) + (1-p)V^B(\alpha_{t+1}(\alpha_t)^-) - pV^B(\alpha_{t+1}(\alpha_t)^+)] > y$$

or, more concisely:

$$\delta_E > \frac{y}{(2p-1)\left[V^B(\alpha_{t+1}(\alpha_t)^+) - V^B(\alpha_{t+1}(\alpha_t)^-)\right]} \tag{48}$$

Note that as $\alpha_t \to 1$, we have that $\alpha_{t+1}(\alpha_t)^- \to \alpha_{t+1}(\alpha_t)^+$ and therefore:

$$\left[V^B(\alpha_{t+1}(\alpha_t)^+) - V^B(\alpha_{t+1}(\alpha_t)^-)\right] \to 0$$

Hence, for any given δ_E, y and p, condition (48) fails to be satisfied.

A.10 Continuous Signals

The expert receives a signal $s \in [0,1]$ about the state of the world which is distributed according to a continuous density function $f_{x,i}(s)$ with distribution function $F_{x,i}(s)$, where $x = 0,1$ and $i = S, D$. For a smart manager (S), the signal structure satisfies the monotone likelihood ratio property (MLRP), that is, $\frac{f_{1,S}(s)}{f_{0,S}(s)}$ is increasing in s. For a dumb manager (D), the signal is uninformative. Without loss of generality, we assume that $f_{1,D}(s) = f_{0,D}(s) = 1$.

Let $\theta \equiv \Pr(x=1)$ be the prior on the state x. Conditional on s we have that

$$\Pr(x=1|s) = \frac{\theta\left[\alpha f_{1,S}(s) + (1-\alpha) f_{1,D}(s)\right]}{\theta\left[\alpha f_{1,S}(s) + (1-\alpha) f_{1,D}(s)\right] + (1-\theta)\left[\alpha f_{0,S}(s) + (1-\alpha) f_{0,D}(s)\right]}$$

Now let $q_1(s) \equiv \alpha f_{1,S}(s) + (1-\alpha) f_{1,D}(s)$ and $q_0(s) \equiv [\alpha f_{0,S}(s) + (1-\alpha) f_{0,D}(s)]$. Hence, we can write:

$$\Pr(x=1|s) = \frac{\theta \frac{q_1(s)}{q_0(s)}}{\theta \frac{q_1(s)}{q_0(s)} + (1-\theta)}$$

Note that $\Pr(x=1|s) > \theta$ if $\frac{q_1(s)}{q_0(s)} > 1$. Given the MLRP, there exists a threshold t^* such that $\Pr(x=1|s) > \theta$ iff $s > t^*$. If the decision maker could observe the signal of the manager, she would update her belief θ upward (downward) whenever $s > t^*$ ($s < t^*$).

Clearly the decision maker does not observe the signal of the expert directly. She only observes a message sent by the expert. Since the decision maker does not receive a private informative signal about the state of the world or the ability of the expert, it is without loss of generality to focus on equilibria in which an expert can send only two messages (that is, all equilibria with more than two messages are payoff-equivalent to the one with only two messages).

In this new setting, we can re-interpret the strategy of truthtelling as the one in which the expert sends $m = 1$ if $s > t^*$, and $m = 0$ otherwise. Provided that DM retains the expert after a correct recommendation and fires her after a mistake, the strategy of truthtelling maximizes the expected probability of being retained. Then, it is straightforward that if the biased expert is sufficiently concerned about the future, there exists an equilibrium in which both biased and unbiased experts report $m = 1$ if $s > t^*$, and $m = 0$ otherwise. This equilibrium corresponds to TT in our setting with a binary signal.

As soon as the biased expert's concern about the future falls below a given threshold, there exists an equilibrium in which the unbiased expert reports truthfully while the biased expert will report $m = 1$ if $s > t_B$, and $m = 0$ otherwise, with $t_B < t^*$ (i.e., the biased expert sends $m = 1$ more often relatively to what prescribed by a truthtelling strategy). This latter equilibrium corresponds to MB.

Now note that the proof that $E_0^{MB}[R_2] > E_0^{TT}[R_2]$ follows the same logic of the proof of proposition 3. Without loss of generality let us assume that $F_{1,S}(s) = F_{0,S}(1-s)$ (this is the equivalent, in the setting with binary signal, of the condition $\Pr(s = 1|S, x = 1) = \Pr(s = 0|S, x = 0) = p$).

In a TT equilibrium, we have that:

$$\begin{aligned}
E_0^{TT}[R_2] = {} & \Pr(0,0,TT)\underbrace{\widehat{\gamma}^{TT}(0,0)}_{\gamma}\underbrace{\widehat{\alpha}^{TT}(0,0,U)}_{\overline{\alpha}} + \Pr(1,1,TT)\widehat{\gamma}^{TT}(1,1)\widehat{\alpha}^{TT}(1,1,U)+ \\
& + [\Pr(0,1,MB) + \Pr(1,0,MB)]\,\gamma\alpha = \\
& = \Pr(0,0,TT)\gamma\overline{\alpha} + \Pr(1,1,TT)\gamma\overline{\alpha} + [\Pr(0,1,TT) + \Pr(1,0,TT)]\,\gamma\alpha = \\
& = \tfrac{1}{2}\Pr(m = 0|x = 0, TT)\gamma\overline{\alpha} + \tfrac{1}{2}\Pr(m = 1|x = 1, TT)\gamma\overline{\alpha}+ \\
& + [\Pr(0,1,TT) + \Pr(1,0,TT)]\,\gamma\alpha
\end{aligned}$$

In an MB equilibrium, we have that:

$$
\begin{aligned}
E_0^{MB}[R_2] = {} & \Pr(0,0,MB)\widehat{\gamma}^{MB}(0,0)\widehat{\alpha}^{MB}(0,0,U) + \Pr(1,1,MB)\widehat{\gamma}^{MB}(1,1)\widehat{\alpha}^{MB}(1,1,U)+ \\
& + [\Pr(0,1,MB) + \Pr(1,0,MB)]\,\gamma\alpha = \\
& = \Pr(0,0,MB)\widehat{\gamma}^{MB}(0,0)\overline{\alpha} + \Pr(1,1,MB)\widehat{\gamma}^{MB}(1,1)\overline{\alpha} \\
& \Pr(0,1,MB)\gamma\alpha + \Pr(1,0,MB)\gamma\alpha = \\
& = \tfrac{1}{2}\Pr(m=0|x=0,MB)\tfrac{\gamma\Pr(m=0|U,x=0,MB)}{\Pr(m=0|x=0,MB)}\overline{\alpha} \\
& +\tfrac{1}{2}\Pr(m=1|x=1,MB)\tfrac{\gamma\Pr(m=1|U,x=1,MB)}{\Pr(m=1|x=1,MB)}\overline{\alpha}+ \\
& + [\Pr(0,1,MB) + \Pr(1,0,MB)]\,\gamma\alpha = \\
& = \tfrac{1}{2}\Pr(m=0|x=0,U,MB)\gamma\overline{\alpha} + \tfrac{1}{2}\Pr(m=1|x=1,U,MB)\gamma\overline{\alpha}+ \\
& + [\Pr(0,1,MB) + \Pr(1,0,MB)]\,\gamma\alpha
\end{aligned}
$$

Note that in MB, U truthtells. Hence, $\Pr(m=0|x=0,U,MB) = \Pr(m=0|x=0,TT)$ and $\Pr(m=1|x=1,U,MB) = \Pr(m=1|x=1,TT)$. This implies that $E_0^{TT}[R_2] - E_0^{MB}[R_2] = [\Pr(0,1,TT) + \Pr(1,0,TT)]\,\gamma\alpha - [\Pr(0,1,MB) + \Pr(1,0,MB)]\,\gamma\alpha$. Intuitively, since in MB the biased expert's signal threshold above which she reports $m=1$ is lower than in TT, the probability of making a mistake is larger in MB than in TT. Hence $E_0^{TT}[R_2] - E_0^{MB}[R_2] < 0$.