Will Truth Out?—An Advisor's Quest To Appear Competent *

Nicolas Klein[†] Tymofiy Mylovanov[‡]

This version: June 26, 2011

Abstract

This paper deals with the problem of providing adequate incentives to an advisor who might be tempted to conceal his true opinion because of his desire to appear competent. We show that if a competent advisor never makes mistakes, the incentive problem will disappear if the time horizon is long enough. If a competent advisor makes infrequent mistakes, the incentive problem will disappear for intermediate time horizons, but will always arise if the time horizon is very long. We furthermore demonstrate that the decision maker can address the incentive problem by letting the advisor accumulate some private information about his ability, and that doing so is optimal if the competent advisor does not make mistakes too often.

Keywords: Reputational cheap talk, career concerns, advisors, strategic information transmission.

JEL Classification Numbers: C73, D83.

---

*We are grateful to Alessandro Bonatti, Matthias Fahn, Johannes Hörner, Navin Kartik, George Mailath, John Morgan, Andrew Postlewaite, Sven Rady, Larry Samuelson, Joel Sobel, and Satoru Takahashi for very useful comments. The first author is particularly grateful to the Department of Economics at Yale University and the Cowles Foundation for Research in Economics for an extended stay, during which this paper took shape. Financial support from the National Research Fund, Luxembourg, and the German Research Fund through SFB TR 15 is gratefully acknowledged.

[†]University of Bonn, Lennéstr. 37, D-53113 Bonn, Germany; email: kleinnic@yahoo.com.

[‡]Department of Economics, University of Pennsylvania; email: mylovanov@gmail.com.

1 Introduction

This paper deals with a decision maker (she) and an advisor motivated by career concerns (he), who interact repeatedly over a finite time horizon. Not fully certain about the quality of his advice, the advisor might decide to play it safe and distort his report toward the decision maker's commonly known prior opinion. The decision maker's goal meanwhile is twofold: She wants to make the best possible use of the current advice she gets, while at the same time learning about her advisor's competence as this will help her make better decisions in the future.

We focus on a situation in which the advisor is either competent or not. Both parties are initially equally uncertain about the advisor's competence. At the end of each period, however, both parties publicly observe if the advisor's prediction has come to pass, and they update their respective opinions about his competence accordingly. Crucially, the decision maker only observes the advisor's report; she does not observe the advisor's information that led him to make the report.

If the advisor were non-strategic and simply told the decision maker whatever he might know, she would employ him until it became clear that he could no longer be of use. Yet, an incentive problem arises because the advisor *is* strategic, and is solely interested in being employed for as long as possible. Thus, he might have incentives to suppress a priori unlikely information in order to maximize his chances of appearing competent; this in turn would slow down the decision maker's learning about his quality and thus render his advice less valuable.

It is well understood that an advisor's concern for appearing competent can create bad myopic incentives, thus decreasing the value of his advice. The literature has so far focused on *single-decision environments* (see, e.g. Trueman (1994), Prendergast and Stole (1996), Scharfstein and Stein (1990), Effinger and Polborn (2001), Levy (2004), Prat (2005), and Ottaviani and Sørensen (2006a, 2006b)).[1]

In contrast, our paper considers a *multi-decision environment.* We demonstrate that in this dynamic environment, forward-looking career concerns create countervailing incentives for the advisor to be truthful. In particular, if a competent advisor never makes mistakes, incentives to report truthfully are restored as the number of periods grows sufficiently large (Proposition

[1] In Morris (2001), the advisor's negative myopic incentives arise because of reputational concerns about his preferences.

3.2). Thus, forward-looking reputational concerns will discipline the advisor's behavior to the point of completely counterbalancing the harmful myopic ones.

Our multi-decision environment gives rise to a novel *bad-news effect*. If the advisor distorts his information in the current period and thereby is successful at appearing competent, he still *privately* gets some negative information about his ability. This in turn makes him more pessimistic about being able to curry favor with the decision maker in future periods. It is thanks to this bad-news effect that, in the case that a competent advisor never makes mistakes, incentives for truthful reporting are restored as the number of periods grows sufficiently large.

If there are myopic incentives to lie, why is it not optimal for the advisor to distort his advice now and postpone truthful reporting until he privately learns that his private information is sound? The key is that distorted reports increase the probability of termination exactly when the advisor is competent and hence decrease the chance of survival into the future periods. Hence, the objectives of appearing competent in the current period and appearing competent over a number of periods are not aligned.

In our analysis, there are two distinct cases. If the competent advisor *never* makes mistakes, the difference between the continuation payoff after the advisor has told the truth and the continuation payoff after he has lied, conditional on the report being correct, is increasing and diverging in the time horizon, making for an unboundedly strong forward-looking reputational effect. Meanwhile, the strength of the myopic reputational effect is invariant in the time horizon; moreover, the advisor's continuation payoff after a mistake is zero as it always results in termination.

Surprisingly, though, if even a competent advisor *occasionally* makes mistakes, the result that truthful reporting is optimal if there are many periods of interaction no longer applies. In fact, the opposite obtains: Incentive problems will always arise if the time horizon is sufficiently long (Proposition 3.3). This is because now the beneficial effects of forward-looking reputational concerns are bounded. The reason for this is that the advisor's payoff from truth-telling now converges as the number of periods increases. Yet, at the same time, the strength of the myopic effect is increasing in the advisor's pessimism about his ability. As the time horizon increases, though, the decision maker is willing to tolerate a larger number of mistakes and, consequently, a more pessimistic advisor. If the time horizon is long enough, there will indeed be a history at which these unfavorable myopic incentives are sufficiently strong to overcome the beneficial forward-looking effect. Thus the bad-news effect is not strong enough to restore truth-telling incentives when the advisor has grown sufficiently pessimistic about his ability.

We reconcile these two different results in Proposition 3.4, where we show for a *fixed time horizon* that if truthful reporting is an equilibrium provided the competent expert never makes mistakes, then it is also an equilibrium if the competent advisor makes mistakes infrequently enough. Thus, the bad-news effect is sufficiently strong: Even if the incentive problem reappears as the time horizon becomes large, the point at which it happens escapes into infinity as the good advisor becomes better informed.

For the case in which the bad-news effect is not strong enough to obviate all incentive problems, and yet the good advisor is still sufficiently competent, we construct the optimal equilibrium and show that the incentive problem is best addressed by letting the advisor gain some private knowledge about his abilities in the first few periods of interaction (Proposition 3.6).[2] Thereafter, the advisor will tell the truth only if he has gained sufficient confidence in his abilities during the previous "grace periods;" otherwise, he will pretend that his information corroborates the common prior perception. This way, the decision maker is only given such information that the advisor, given his superior private information, deems valuable enough; his white lies, on the other hand, are inconsequential, in the sense that a decision maker who knew what he knew would ignore this information also. Moreover, putting up with the advisor's occasional white lies avoids the decision maker the cost of sometimes losing the valuable services of an advisor whose only fault has been correctly to predict the expected.

Our analysis has implications for how best to capture the bad-news effect in a reduced-form model of career concerns with one period of interaction. In particular, it is important to allow the advisor's utility increment from reputation to depend both on the public belief and on his *private belief* about his competence. By contrast, the existing models often assume that the value of reputation is fully determined by the public belief. The results in this paper suggest that it can be appropriate to assume that the value of reputation is increasing in the private belief and that this effect can counteract the bad reputation incentive caused by the desire to improve the public belief.

In our paper, communication is cheap talk. The seminal paper in this literature is Crawford and Sobel (1982). The early references on cheap talk with reputational concerns about

[2] Endogenous accumulation of private information, albeit off the equilibrium path, can also occur in Bergemann and Hege (1998, 2005) and Hörner and Samuelson (2009), who examine a dynamic agency problem in which an agent can conceal funds and divert them toward his private ends.

preferences are Sobel (1985) and Benabou and Laroque (1992). Negative reputational effects due to preference uncertainty appear in Morris (2001) and Ely and Välimäki (2003). The structure of the incentive problem is quite different in these models, in particular because the advisor knows his preferences and hence the bad-news effect is not relevant.

Career concerns for expertise are studied in Trueman (1994), Prendergast and Stole (1996), Scharfstein and Stein (1990), Effinger and Polborn (2001), Suurmond, Swank, and Visser (2004), Levy (2004), Prat (2005), and Ottaviani and Sørensen (2006a, 2006b).[3] While there is a single advisor in some of these models, other papers consider multiple experts and focus on incentives for herding or for contrarian reports. Dasgupta and Prat (2006, 2008) show that financial traders' career concerns and their pursuit of a reputation for expertise can increase trading volumes and prevent asset prices from reflecting fundamental values. Levy (2007) looks at career concerns in committee decision-making.[4]

There is a more distant connection between our paper and the literature on the testing of experts (see e.g., Foster and Vohra (1998, 1999), Olszewski and Sandroni (2008), and Shmaya (2008)). The paper most closely connected to ours in this literature is Olszewski and Peski's (forthcoming) infinite horizon principal-agent model. In this literature, experts privately know their type and the objective is to construct a test separating experts of different types. By contrast, the objective of the decision maker in our model is to induce the advisor to report his private signals truthfully. The advisor does not know his type and the optimal decision rule does not always separate different types.

Our investigation is also related to Holmström's (1999) seminal contribution on career concerns and the subsequent related literature. However, in contrast to Holmström (1999), our advisor's career concerns reveal themselves through his cheap-talk communication rather than his choice of costly effort.

The rest of this paper is structured as follows: Section 2 presents the model, and introduces necessary notation; in Section 3, we analyze the first-best and the second-best decision rules; Section 4 concludes. The proofs omitted in the main text and extensions are provided in the Appendix.

[3]See also Morgan and Stocken (2003) who analyze a model with uncertainty about the expert's relative preference between inflating his reports and providing an accurate forecast.

[4]Bourjade and Jullien (2004) offer a model of expertise with reputational concerns with hard information.

2 Model Setup

We study the simplest model that formalizes the advisor's reputational concerns in an explicitly dynamic setting. In our model, there are $N \geq 2$ periods. In each period, a decision maker chooses a policy. The optimal policy is uncertain and is described by the random variable $\omega_t \in \{0, 1\}$, which is iid across periods and is equal to 1 with a commonly known probability $p \in (0, 1/2)$. In each period, the decision maker's payoff is 1 if the policy matches the state and 0 otherwise; it is publicly revealed at the end of the period. There is no discounting.

The decision maker can consult an advisor before making a policy choice. The advisor does not care about the decision maker's policy choices; his only objective is to be consulted as often as possible. Specifically, he gets a payoff of 1 per period when he is employed and 0 otherwise. Again, there is no discounting.

If consulted, the advisor first observes a binary noisy non-verifiable signal $\tilde{s} \in \{0, 1\}$ about the realization of the state; then he sends a cheap-talk message to the decision maker about what he has observed. The quality of the signal is initially unknown and believed by both parties to be high with probability $\alpha \in (0, 1)$ and low with the counter-probability. The low-quality signal is uninformative and is always equally likely to be correct or incorrect, whichever the realized state may be. The high-quality signal is informative and is correct with a time-invariant commonly known probability $q \in (1 - p, 1]$. The signals are iid across periods. We refer to the quality of the signal as the advisor's *competence*.

We denote by α_t the decision maker's belief about the advisor's competence at the beginning of period t; we refer to it as the advisor's *reputation*. The advisor's corresponding belief is denoted by $\hat{\alpha}_t$. This belief could well differ from the decision maker's because the advisor has the benefit of privately knowing the signals he has observed.

To rule out uninteresting cases, we impose the following

Assumption 2.1 *It is commonly believed that* $\alpha q + (1 - \alpha)/2 < 1 - p$;

i.e. the decision maker obtains a higher payoff if she follows her prior beliefs than if she follows the signals of an advisor with reputation α. Simultaneously, this assumption implies that an advisor with a reputation of α will believe that state 0 is more likely regardless of his signal and hence he might have incentives to lie about his signal.

The timing of the interaction in each period is as follows. First, the decision maker decides whether to hire the advisor. If he is employed, the advisor then observes a signal and sends a subsequent cheap-talk report to the decision maker, after which the decision maker chooses a policy. Then, at the end of the period, the actual state of the world is publicly observed, and payoffs are realized. Our solution concept is perfect Bayesian equilibrium.

In order to focus on the advisor's incentives and to clarify the core intuition behind our main insights, we restrict the decision maker's behavior and require that she terminate the advisor if there is no value in continuing to employ him.

Assumption 2.2 *We restrict attention to those equilibria in which the decision maker terminates the advisor whenever the benefit of continuing to employ him is 0.*

This restriction could be viewed as a reduced-form representation of behavior in a richer model in which the decision maker has limited commitment power and incurs an opportunity cost of employing the advisor. This cost could e.g. represent exogenously specified wages, opportunity costs of the decision maker's time spent with the advisor, or resources required to provide the advisor with access to information. In some applications, this restriction could also be a consequence of external political pressures that make it impossible to retain an advisor who has proved himself to be incompetent. Indeed, without Assumption 2.2, the advisor's career concerns would have no impact in our model because it would be optimal for the decision maker simply never to fire the advisor.

Our core intuition that the bad-news effect counteracts the advisor's myopic incentives to lie extends to richer environments that assume time-invariant positive costs of employment or the advisor's being paid his value-added for the decision maker as his salary, as we show in the appendix.

In our environment, the decision maker faces two objectives. On the one hand, she chooses an optimal policy in each period given the available information. On the other hand, she chooses her employment strategy with a view toward minimizing the effect the advisor's career concerns will have on his reports. Achieving the first objective is straightforward and will not be the focus of our analysis: If the advisor is employed, the decision maker will follow his recommendation if and only if it is sufficiently informative in expectation. In particular, if the decision maker believes the advisor is telling the truth, following his report is strictly optimal if and only if the decision maker thinks the signal is informative enough to overcome her prior, i.e.

$$\alpha_t q + (1 - \alpha_t)/2 > 1 - p. \tag{1}$$

If the advisor's report is not sufficiently informative or if the advisor is not consulted, the decision maker will follow her prior and choose policy 0. Assumption 1 states that (1) does not hold in the first period; hence, the decision maker will always implement policy 0 in the first period.

3 Optimal Decision Rules

As our first-best benchmark, we consider a hypothetical environment in which the advisor's signals are observed by the decision maker.[5] Let $\alpha_N(k)$ denote the posterior belief that the advisor is competent at the beginning of the last period if there were k incorrect signals in the preceding periods. The value of $\alpha_N(k)$ is positive and decreasing in k if $q < 1$ and is equal 0 for any $k \geq 1$ if $q = 1$. The advisor's signal in the last period is valuable for the decision maker if following the signal generates a higher expected payoff than following her prior, i.e. if

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

To avoid uninteresting cases, we make

Assumption 3.1 *The inequality (2) is satisfied for $k = 0$.*

Definition Let κ be the highest $k \in \mathbb{N} \cup \{0\}$ for which (2) is satisfied.

Thus, κ is the maximal number of mistakes after which the advisor's signal is valuable for the decision maker in the last period.

Definition The first-best decision rule

1. employs the advisor until his reports have disagreed with the state $\kappa + 1$ times;

2. implements a policy equal to the advisor's report if $\alpha_t > \frac{1-2p}{2q-1}$ and policy 0 otherwise.

[5] Alternatively, we could think of an advisor who has no career concerns and is committed to report his signals truthfully.

If the advisor's reports are truthful, this rule is a best response for the decision maker because it maximizes her payoff and retains the advisor if and only if the decision maker's continuation value from doing so is positive. The first-best decision rule provides a natural benchmark against which to assess the effect of the advisor's career concerns. Furthermore, the decision maker's payoff if she follows the first-best decision rule and the advisor reports his signals truthfully is the upper bound on her payoff in our model as well as in a richer model in which consulting an advisor entails a small opportunity cost (cf. our remarks after Assumption 2.2).

Definition The first-best decision rule is *incentive compatible* if there exists an equilibrium in which the decision maker follows this rule and the advisor's reports are truthful for every history on the equilibrium path.

The agency problem in our model arises because the first-best decision rule might not be incentive compatible. Let, for instance, $N = 2$ and $\kappa = 0$, and imagine that the advisor observes $\tilde{s}_1 = 1$ in the first period. By Assumption 2.1, condition (1) is violated with slackness for $t = 1$ and, therefore, the advisor believes that the state $\omega_1 = 0$ is more likely. Thus, if the decision maker followed the first-best rule, the expert would maximize his probability of employment in the next period by reporting $\hat{s}_1 = 0$. As a result, the advisor's best response to the first-best decision rule would entail a report of 0 in period 1 irrespective of the observed signal.

If a competent advisor never makes mistakes, the following proposition shows that if N exceeds a certain threshold, the first-best decision rule becomes incentive compatible. The result relies on what we call the *bad news* effect: If the advisor lies and his report turns out to be correct, he privately learns that he is incompetent. By contrast, if he reports his signal truthfully and it is correct, then the advisor believes that he is more likely to be competent. Moreover, if the report is incorrect, the advisor is fired and his continuation payoff is 0 regardless of his beliefs. Although there is an obvious analogy, the proof is not a folk-theorem type of argument. First of all, there is no discounting in our environment and the number of periods is finite. More importantly, the incentive problem disappears because of the different rates of growth in the payoffs from lying and from telling the truth as the number of periods increases. The reason for this is that the advisor evaluates his future payoffs conditional on different events.

Proposition 3.2 (Vanishing Career Concerns) *Assume that the competent advisor never makes mistakes. For any given p and α, there exists an integer $\check{N}_0$ such that the first-best decision rule is incentive compatible if and only if $N \geq \check{N}_0$.*

PROOF: A formal version of the argument expounded above proves that for any t there exists an integer $N'(t)$ such that for all $N \geq N'(t)$ there is no profitable (possibly, multi-period) deviation from truth-telling that starts in period t. It is left to show, then, that there exists an $\check{N}_0$ such that $N'(t) \leq \check{N}_0$ for all t or, in other words, that as we increase N the incentive constraints are not violated in the newly added periods. This, however, holds true because, if the advisor is employed toward the end of the relationship under the first-best decision rule, then his reputation is necessarily high, the advisor considers his signals very informative, and truth-telling is his strict best response. A complete proof is provided in the appendix. ∎

The insight that a longer time horizon solves the incentive problem is valid if the competent advisor is always correct. However, if the competent advisor might occasionally observe incorrect signals, this is no longer the case, as the following example shows. Here, the first-best outcome can be attained in equilibrium if $N = 2$ but not if $N = 3$.

Example Let $\alpha = 5/12$, $p = 3/7$, and $q = 9/10$.

1. Let $N = 2$. The first-best decision rule retains the advisor in period 2 if and only if his signal is correct in period 1. This rule is incentive compatible.

2. Let $N = 3$. The first-best decision rule always retains the advisor in period 2, and retains him in period 3 if and only if his signal was correct at least once in the previous two periods. This rule is *not* incentive compatible. In particular, if the decision maker follows this rule, the advisor's best response after an incorrect signal in period 1 is to disregard his signal and report 0 in period 2.

In this example, the decision maker would like to continue to employ the advisor if he makes a mistake in period 1 if $N = 3$ but not if $N = 2$. This is so because with more remaining periods there is a chance that the advisor will prove himself to be sufficiently competent to become valuable for the decision maker. However, after a mistake, the advisor is no longer willing to report his signal truthfully. If $N = 2$, this does not matter as the advisor is fired but if $N = 3$ the first-best decision rule ceases to be incentive compatible. This difficulty does not arise if $q = 1$, as then a single mistake fully reveals that the advisor is of no value to the decision maker. It is true in general that if $q < 1$ there will arise a history at which the advisor is too pessimistic to tell the truth if the time horizon is long enough, as the following proposition shows:

Proposition 3.3 (Persistent Career Concerns) *Suppose the competent advisor occasionally makes mistakes, i.e. $q < 1$. For any given p and α, there exists an integer N_0 such that the first-best decision rule is* not *incentive compatible if $N \geq N_0$.*

PROOF: See Appendix. ∎

This surprising discontinuity notwithstanding, our next proposition shows that the problem still behaves continuously in the following sense: As the probability of a competent advisor's making a mistake vanishes, the number of periods of interaction needed to make the first best incentive incompatible will diverge. The reason is that the advisor's incentives are continuous in the competent type's probability of being correct q.

Proposition 3.4 *For any α and p, there exists $q_0 \in (1-p, 1)$ such that the first-best decision rule is incentive compatible if $q \geq q_0$ and $\underline{N}(q) \leq N \leq \overline{N}(q)$ for some $\underline{N}(q), \overline{N}(q)$, where $\check{N}_0 \leq \underline{N}(q) \leq \overline{N}(q)$. Furthermore, $\underline{N}(q) - \check{N}_0 \leq 1$ for $q \geq q_0$ and $\overline{N}(q) \to \infty$ as $q \to 1$.*

PROOF: See Appendix. ∎

We now turn to environments in which the first best is not incentive compatible. A quite natural way for the decision maker to handle the advisor's incentive problem would be for her to grant him an initial "grace stage," during which he was allowed to send uninformative signals each period, and to gain confidence in his abilities, finding his mark in his new job. Once this probationary phase ends, though, he is expected to be right every time, i.e. he is fired as soon as he makes a mistake. The advisor will then report his signals truthfully if his signals have all been correct during the probationary phase; otherwise, he may well best respond by continuing to babble, i.e. to announce state 0 no matter what his signal may have been.

We summarize this equilibrium in the next proposition. In order to do so, we first define the period t^{FB} as follows: Assume the decision maker follows the first-best decision policy. Now, let t^{FB} be the earliest period such that an advisor who has observed and reported only correct signals, including in this period, will henceforth find truthful reporting optimal.[6] (Clearly, if $t^{FB} = 0$, the first-best decision rule is incentive compatible. Furthermore, $t^{FB} < N$ because the advisor is indifferent about his report in the last period.)

Proposition 3.5 (Equilibrium With A Grace Stage) *There exists an equilibrium in which no information is transmitted, and the advisor is never fired, during the first t^{FB} periods; thereafter, the advisor truthfully reveals his signals if his first t^{FB} signals were correct. Moreover, he will only be fired as soon as he has made an incorrect forecast after the first t^{FB} periods.*

Proof: Let τ be the current period. Now, the advisor's equilibrium strategy is specified as follows: (0) If he has reported 1 in one or more of the first t^{FB} periods or made an incorrect report in a period in $\{t^{FB}+1, \cdots, \tau-1\}$, he will report 0 in period τ. After those histories that are not covered by statement (0), the advisor will (i) report 0 in all periods $\tau \leq t^{FB}$; (ii) will report his signals truthfully if $\tau > t^{FB}$ and all of his signals in the first t^{FB} periods were correct; (iii) if $\tau > t^{FB}$ and he has observed an incorrect signal in the first t^{FB} periods, he will report the state that seems more likely to him given his signal.[7]

The decision maker's equilibrium strategy calls for not hiring the advisor in those periods τ such that there exists a period $\tilde{\tau} < \tau$ in which the advisor has given an incorrect forecast and $\tilde{\tau} > t^{FB}$, or in which the advisor has reported 1 and $\tilde{\tau} \leq t^{FB}$. In all other periods, she employs the advisor.

These strategies are mutually best responses by the definition of t^{FB}. ∎

Now, let us consider the case of $\kappa = 0$. The decision maker's policy choices in this equilibrium are those she would make in the first-best environment: In each period during the grace phase, the decision maker implements policy 0. She would take the same action in the first-best environment because she is still pessimistic about the quality of the expert's signal. After the grace phase, a report of 1 reveals that the expert is truthful and has only observed correct signals thus far, which allows the decision maker to take the first-best action. The report of 0 does not reveal the private history of the advisor; this is inconsequential, however,

[6] That is, the advisor's optimal strategy in period $t^{FB}+1$ prescribes truthful reporting in this period and in each period $t > t^{FB}+1$ provided the report in periods $t^{FB}+1, \cdots, t-1$ were also truthful. The history of actual signal realizations either before or after t^{FB} is immaterial here.

[7] If $\kappa = 0$, this always implies babbling, i.e. reporting state 0.

as action 0 is the decision maker's best response even if the advisor had been fired in the first-best environment. At the optimum, the first-best quality of policy decisions is thus achieved thanks to a longer ex-ante expected duration of employment than in the first best.

Indeed, in our model, it can only be to the principal's advantage for the agent to be better informed, even if this information be held privately; an advisor who is more optimistic will be more inclined to reveal his signal, and following his signal is a good idea for the principal also. A privately pessimistic advisor by contrast will tend to report his prior without any regard to his signal; in this case, following her prior belief is also the best the principal can do in terms of policy. If, on the other hand, the principal's primary goal were to screen out a bad advisor, private information would rather tend to hurt the principal.[8]

Thus, even though the first-best decision rule may not be incentive compatible, this equilibrium still achieves the first-best payoff for the decision maker that she would attain in the environment in which the advisor's information is public. Nevertheless, if $t^{FB} \geq 1$, the equilibrium violates condition 1. of our definition of the first best, as the advisor is employed longer in expectation than in the first-best rule (recall from our discussion after Assumption 2.2 that our model could be viewed as a reduced-form representation of an environment in which consulting an advisor entails a small cost for the decision maker). Of course, if the decision maker incurred such a (small) cost for employing the advisor, she would prefer firing a bad advisor as quickly as possible. As it turns out, it is impossible to achieve the first-best payoff while employing the advisor for fewer expected periods than in our equilibrium, as the following proposition shows. Thus, this equilibrium would continue to be second-best in a richer model with employment costs, provided these costs were sufficiently small.

Proposition 3.6 (Second-Best Optimum) *If $\kappa = 0$, the decision maker's ex-ante expected payoff in the equilibrium identified in Proposition 3.5 is equal to the first-best payoff. Furthermore, there does not exist an equilibrium in which the decision maker obtains the same ex-ante expected payoff and the ex-ante expected duration of the advisor's employment is lower.*

Proof: The first statement immediately follows from our previous discussion. Regarding the second statement, suppose on the contrary that there exists an equilibrium achieving the first-

[8] In Olszewski and Peski (forthcoming), the first best is also approached thanks to a "grace stage," which performs quite a different function in their model: As their advisor is already perfectly informed about his type, there is no need for him to accumulate private information, and hence he will not simply be babbling during his grace stage.

best payoff in which the advisor is employed for fewer periods in expectation. In order for the principal to achieve an ex-ante expected value of the first-best payoff, it must be the case that a good advisor is never fired; i.e. in such an equilibrium, the advisor is only fired after he has revealed himself to be of the bad type. Since he is employed for fewer periods in expectation than in the equilibrium exhibited in Proposition 3.5, it must be the case that some information on the agent's type will be transmitted in period t^{FB} or earlier. If $t^{FB} = 0$, this is impossible. If $t^{FB} \geq 1$, the decision maker has to fire the advisor with some positive probability even after he has been correct, in order to induce him to tell the truth with some positive probability in period t^{FB} or earlier, since Assumption 2.2 rules out keeping the advisor on after he has made a mistake. This in turn implies that a good advisor will be fired with positive probability. Hence, the decision maker makes worse policy decisions in expectation, and thus her payoff is bounded away from the first-best payoff. ∎

If $\kappa > 0$, the characterization of the second-best optimal equilibrium becomes much more involved. The basic insight, though, that allowing the agent to accumulate some private information about his type might help alleviate incentive problems is not particular to the case of $\kappa = 0$. However, the principal might now avail herself of many different ways of allowing the agent to accumulate this private information; e.g. there may well be a sequence of nonconsecutive blocks of grace periods, with the agent being moved back into such a block of appropriate length after he has made a mistake in a phase of play in which he was expected to tell the truth. Also, the first grace period need no longer coincide with the first period of play. We leave a rigorous exploration of these issues outside the scope of this paper.

4 Conclusion

We have investigated the dynamic interaction between a decision maker and an advisor of unknown quality who privately observes a potentially decision-relevant signal. As he only cares about his reputation insofar as it translates into a longer expected duration of employment, the advisor may have incentives strategically to manipulate the cheap-talk relay of his signal to the decision maker. We have shown that if a competent advisor never makes mistakes and the number of periods is large enough, the impact of the advisor's career concerns vanishes, and the first best becomes implementable; however, the opposite is true if a competent advisor occasionally makes mistakes. Moreover, we have shown that the decision maker can address the incentive problem by letting the advisor accumulate some private information about his ability; doing so is optimal if a competent advisor only makes mistakes very infrequently.

In our model, the decision maker can only set incentives by either retaining or firing the advisor. In this setting, we have seen that encouraging inconsequential chatter can be the optimal way to proceed. However, in some economic situations, the decision maker might be in a position to hide the realization of the actual state from the advisor. We would conjecture that our decision maker would want to do so if she was faced with an optimistic advisor, thus shielding him from potentially bad news, which might make him coyer about revealing his signals in the future. Whereas she might thus be able to slow down the advisor's learning about his type, she would not be able completely to shut it down, as the advisor could still draw inferences about his type from the relative frequency of the different signal realizations. By contrast, the decision maker would want to reveal the outcomes of her policy to pessimistic advisors, so as to expedite their learning process. We leave a full exploration of these issues for future work.

Appendix

A Proofs

Proof of Proposition 3.2

Suppose the decision maker pursues the first-best policy of immediately firing the advisor if, and only if, the advisor has made a mistake. Then, the agent is willing to reveal a signal indicating the less likely state 1 truthfully at any time t, if at all times $1 \leq t \leq N$, the following incentive constraint holds:

$$p\left[\alpha_t(N-t) + \frac{1-\alpha_t}{2}\left(1+\tfrac{1}{2}+\cdots+\frac{1}{2^{N-t-1}}\right)\right] \\ \geq (1-p)\frac{1-\alpha_t}{2}\left[1+(1-p)+\cdots+(1-p)^{N-t-1}\right], \quad \text{(A.1)}$$

where α_t is the posterior belief about the advisor's competence provided all his signals have been correct. To understand the right-hand side of the incentive constraint, the reader should note that if, upon lying, the advisor finds out *ex post* that his message was in fact correct, he then privately learns that he is of the low type and will maximize his continuation payoff by reporting the *a priori* more likely state in all subsequent periods.

It is now immediate to verify that, as $N \to \infty$, the left-hand side diverges to $+\infty$, whereas the right-hand side converges to $\frac{1-p}{p}\frac{1-\alpha_t}{2} < \infty$. Let $\check{N}_0$ be the smallest value of N for which this constraint is satisfied for all $t \leq K$, where we define $K := \log_2\left(\frac{1-2p}{p}\frac{1-\alpha}{\alpha}\right)$. By our Assumption 2.1, we have that $\check{N}_0 \geq 2$.

Let $N = \check{N}_0$. Then, the constraint is also satisfied for all $t > K$: It is immediate to verify that the constraint holds for any N if $\alpha_t = 1 - 2p$. Furthermore, the left-hand side of the constraint is increasing in α_t while the right-hand side is decreasing in α_t. Therefore, the constraint is satisfied for all $\alpha_t \geq 1 - 2p$, which is equivalent to $t \geq K$.

As is straightforward to verify, the left-hand side of the incentive constraint conditional on a signal indicating the more likely state 0, is $\frac{1-p}{p} > 1$ times the left-hand side of the above constraint, whereas the right-hand side is $\frac{p}{1-p}$ times the above right-hand side. Therefore, the incentive constraint after signal 0 also holds for $N = \check{N}_0$.

To complete the proof, it is sufficient to show that (A.1) holds for all $N > \check{N}_0$. Let $H(N,t)$ be the slack in the incentive constraint (A.1), i.e., the difference between the left-hand side and the right-hand side of the constraint. Then,

$$\mathrm{sign}[H(N+2,t) - H(N+1,t)] - [H(N+1,t) - H(N,t)] \quad = \quad \mathrm{sign}\left[(1-p)^{N+1-t} - \frac{1}{2^{N+1-t}}\right]$$

and hence $H(N)$ is discretely strictly convex (Yüceer 2002). Therefore, since by Assumption 2.1 $H(1) < 0$ we have that if $H(\check{N}_0) \geq 0$, then $H(N) > 0$ for $N > \check{N}_0$.

Proof of Proposition 3.3

Fix arbitrary parameters α, p and $q < 1$. Let h^* be a history such that (1) the advisor has always reported truthfully, (2) all of his reports have been incorrect, and (3) one additional incorrect report will result in termination of employment. A necessary condition for the first-best decision rule to be incentive compatible is that a deviation from truthfully reporting a signal of 1 to reporting 0 in the current period and all future periods not be profitable at history h^*. Let α' be the advisor's belief about his competence, and $k = N - t$ the remaining number of periods at h^*. Then, this condition can be expressed as

$$\begin{aligned} p\left[\alpha'\left(q + q^2 + \cdots + q^k\right) + (1-\alpha')\left(\frac{1}{2} + \frac{1}{4} + \cdots + \frac{1}{2^k}\right)\right] \\ \geq (1-p)\left[(1-\alpha')\frac{1}{2} + \alpha'(1-q)\right]\left[1 + (1-p) + \cdots + (1-p)^{k-1}\right], \quad \text{(A.2)} \end{aligned}$$

or, equivalently,

$$\begin{aligned} \alpha'\left(p\frac{q}{1-q}(1-q^k) - p\left(1 - \tfrac{1}{2}^k\right) + (1-p)(q - \tfrac{1}{2})\frac{1-(1-p)^k}{p}\right) \\ \geq (1-p)\frac{1-(1-p)^k}{2p} - p\left(1 - \tfrac{1}{2}^k\right), \quad \text{(A.3)} \end{aligned}$$

The left-hand side is increasing in k and converges to $\alpha'(2q-1)\left(\frac{p}{1-q}+\frac{1-p}{2p}\right)$ from below, while the right-hand side is also increasing in k and converges to $\frac{1-p}{2p}-p$ from below. Therefore, if

$$\alpha' < \alpha^* := \frac{(1-p-2p^2)(1-q)}{(2q-1)(2p^2+1-p-q+pq)},$$

there exists K^* such that for all $k \geq K^*$, (A.2) is violated.

To prove the statement of the proposition, we need to establish that as N diverges, both κ and $N-\kappa$ diverge. Indeed, if κ diverges then the advisor's belief about his competence at h^* converges to 0 and will be below α^* if N is sufficiently large. If, in addition, the number of remaining periods at history h^*, which is $N-\kappa$, diverges, then there exists N_0 such that (A.2) is violated for all $N \geq N_0$.

The value of κ is the largest integer k that satisfies:

$$\left(\frac{1-q}{q}\right)^k > \frac{1-\alpha}{\alpha}\frac{\frac{1}{2}-p}{q-(1-p)}\left(\frac{1}{2q}\right)^{N-1}. \tag{A.4}$$

From (A.4), we have that as N diverges, the right-hand side converges to 0 and hence κ diverges. At the same time, (A.4) can be rewritten as

$$\left(\frac{q}{1-q}\right)^{N-\kappa} > \frac{1-\alpha}{\alpha}\frac{\frac{1}{2}-p}{q-(1-p)}\frac{q}{1-q}\left(\frac{1}{2(1-q)}\right)^{N-1}.$$

The right-hand side diverges in N and hence $N-\kappa$ diverges. ∎

Proof of Proposition 3.4

Let $\overline{N}_1(q)$ be the largest N such that $\kappa = 0$ or, equivalently,

$$\frac{1-q}{q} \leq \frac{1-\alpha}{\alpha}\frac{\frac{1}{2}-p}{q-(1-p)}\left(\frac{1}{2q}\right)^{N-1}.$$

Note that $\overline{N}_1(q) \to \infty$ as $q \to 1$.

Two necessary conditions for the first-best decision rule to be incentive compatible is that, at $t=1$ and $t=N-1$, a deviation from truthfully reporting a signal of 1 to reporting 0 not be profitable. These conditions can be expressed as

$$\begin{aligned} p\left[\alpha\left(q+q^2+\cdots+q^{N-1}\right)+(1-\alpha)\left(\frac{1}{2}+\frac{1}{4}+\cdots+\frac{1}{2^{N-1}}\right)\right] \\ \geq (1-p)\left[(1-\alpha)\frac{1}{2}+\alpha(1-q)\right]\left[1+(1-p)+\cdots+(1-p)^{N-2}\right], \end{aligned} \tag{A.5}$$

and

$$\alpha q^{N-1}+(1-\alpha)\frac{1}{2^{N-1}} \geq (1-p)\left(\alpha q^{N-2}+(1-\alpha)\frac{1}{2^{N-2}}\right). \tag{A.6}$$

From the proof of Proposition 3.2, we know that both conditions hold with slackness for $q=1$ if $N > \check{N}_0$. Moreover, observe that the left-hand side and the right-hand side of both conditions are continuous in q and N. Therefore, for a sufficiently high value of q, there exists an integer N for which these conditions are satisfied. Let $\underline{N}(q)$ be the smallest such integer. The value of $\underline{N}(q)$ is non-increasing in q.

We shall now show that conditions (A.5) and (A.6) are also sufficient for incentive compatibility if $N \leq \overline{N}_1(q)$: To see this, consider an advisor's strategy s that calls for (a) reporting truthfully in the first period and in all subsequent periods if the advisor has always reported truthfully in the past and all his reports were correct, and (b) for reporting 0 otherwise. By our definition of $\overline{N}_1(q)$, the decision maker's first-best strategy is to retain the advisor if and only if all of his previous reports were correct.

We have to consider two kinds of one-shot deviations for the advisor. First, the advisor can deviate and report 0 after a signal of 1 on the equilibrium path at a history in which all of his previous reports were truthful and correct. Second, he can deviate by reporting a signal of 1 off the equilibrium path at a history in which one of his previous reports was not truthful but correct.

Consider the first kind of deviation and assume it happens in period $t+1$. Note that we can safely ignore deviations in period N, since the expert is indifferent over what to report in the last period. Hence, $t \in \{0, \cdots, N-2\}$. The advisor's ex-ante expected payoff under this deviation equals

$$\begin{aligned} U'_t = 1 + \alpha \left[q + \cdots + q^t + q^t(1-p) + \cdots + q^t(1-p)^{N-t-1}\right] \\ + (1-\alpha)\left[\frac{1}{2} + \cdots + \frac{1}{2^t} + \frac{1}{2^t}(1-p) + \cdots + \frac{1}{2^t}(1-p)^{N-t-1}\right] \end{aligned}$$

if $t \in \{1, \cdots, N-2\}$. For $t = 0$, i.e. a deviation in the first period, we have

$$U'_0 = 1 + (1-p) + \cdots + (1-p)^{N-1}.$$

A simple calculation gives

$$U'_t - U'_{t+1} = \left[\alpha q^t((1-p)-q) + (1-\alpha)\frac{1}{2^t}\left((1-p) - \frac{1}{2}\right)\right]\left[1 + (1-p) + \cdots + (1-p)^{N-t-2}\right]$$

for all $t \in \{0, \cdots, N-3\}$. For $t = N-2$, i.e. a deviation in the second to last period, we have that

$$U'_{N-2} - U'_{N-1} = \alpha q^t((1-p)-q) + (1-\alpha)\frac{1}{2^t}\left((1-p) - \frac{1}{2}\right).$$

Thus, for all $t \in \{0, \cdots, N-2\}$, we have that $U'_t - U'_{t+1} \geq 0$ if, and only if

$$\alpha q^t((1-p)-q) + (1-\alpha)\frac{1}{2^t}\left((1-p) - \frac{1}{2}\right) = \alpha_{t+1}(k=0)q + (1-\alpha_{t+1}(k=0))\tfrac{1}{2} \leq 1-p.$$

As $\alpha_t(k=0)$, and hence $\alpha_t(k=0)q + (1-\alpha_t(k=0))\frac{1}{2}$, is increasing in t, we have that: (i) If $U'_t - U'_{t+1} \geq 0$, then $U'_{t'} - U'_{t'+1} > 0$ for all $t' < t$; and (ii) if $U'_t - U'_{t+1} \leq 0$, then $U'_{t'} - U'_{t'+1} < 0$ for all $t' > t$. It thus follows that conditions (A.5) and (A.6) are also sufficient to deter deviations of the first kind.

To prove that the advisor cannot profit from a deviation of the second kind, i.e. off the equilibrium path, let α' denote the advisor's private belief about his competence and $\tilde{K}$ the number of remaining periods. The advisor finds it optimal to report 0 after a signal of 1 if

$$\left(\alpha' q + (1-\alpha')\frac{1}{2}\right)\left[(1-p) + (1-p)^2 + \cdots + (1-p)^{\tilde{K}-1}\right] \leq (1-p)\left[(1-p) + (1-p)^2 + \cdots + (1-p)^{\tilde{K}-1}\right] \quad \text{(A.7)}$$

Observe that the advisor can reach this history only if he has deviated on the equilibrium path and made an untruthful report that turned out to be correct. Then, by our definition of $\overline{N}_1(q)$, α' is sufficiently small for the constraint to be satisfied.

Thus, we have shown that our conditions (A.5) and (A.6) are also sufficient for incentive compatibility. By continuity of (A.5) and (A.6) in q and N, we can hence conclude that $\underline{N}(q) \geq \check{N}_0$, and that there exists a $q_1 < 1$ such that $0 \leq \underline{N}(q) - \check{N}_0 \leq 1$ for all $q \geq q_1$.

It is direct to verify that (A.6) is also satisfied for all $N \geq \underline{N}(q)$. We now show that

(*) there exists $\overline{N}_2(q)$, with $\overline{N}_2(q) \to \infty$ as $q \to 1$ such that (A.5) is also satisfied if $\underline{N}(q) \leq N \leq \overline{N}_2(q)$.

Let $F(N)$ be the difference between the left-hand side and the right-hand side of (A.5). Then,

$$\operatorname{sign}[F(N+2,t) - F(N+1,t)] - [F(N+1,t) - F(N,t)] = = \operatorname{sign}\left[-\alpha(1-q)\left[q^N - (1-p)^N\right] + \frac{1-\alpha}{2}\left[(1-p)^N - \frac{1}{2^N}\right]\right]. \quad \text{(A.8)}$$

Let $\overline{N}_2(q)$ be the largest integer such that this sign is positive for all $N \leq \overline{N}_2(q)$ and hence $F(N)$ is discretely strictly convex (Yüceer 2002) for $N \in \{1, \ldots, \overline{N}_2(q)\}$. Clearly, as $q \to 1$, we have $\overline{N}_2(q) \to \infty$. To see this, observe that for any N, there exists $q_N < 1$ such that $\alpha(1-q)\left[q^N - (1-p)^N\right] < \frac{1-\alpha}{2}\left[(1-p)^N - \frac{1}{2^N}\right]$ for all $q > q_N$. Now, for any N', choose $q^*_{N'} = \max_{N \leq N'} q_N$. Then, if $q > q^*_{N'}$ we have that the sign in (A.8) is positive for all $N \leq N'$. This implies that if $q > q^*_{N'}$, then $\overline{N}_2(q) \geq N'$. Since the choice of N' is arbitrary, the argument is complete. Therefore, (*) holds and there exists some $q_2 \in (1-p, 1)$ such that $\check{N}_0 + 1 \leq \overline{N}_2(q)$ for all $q > q_2$.

Now, set $\overline{N}(q) = \min\left\{\overline{N}_1(q), \overline{N}_2(q)\right\}$ and choose q_0 such that $1 > q_0 \geq \max\{q_1, q_2\}$ and $\overline{N}(q) \geq \check{N}_0 + 1$ for all $q \geq q_0$. Since $\overline{N}(q) \to \infty$ as $q \to 1$, such a $q_0 < 1$ exists. Then, by construction, $\underline{N}(q) \leq \overline{N}(q)$, and the first best is incentive compatible for all $N \in \{\underline{N}(q), \cdots, \overline{N}(q)\}$, for all $q \geq q_0$. ∎

B Extensions

Costly advice

Consider a modified model in which the decision maker incurs an exogenously specified cost $c \in (0, p)$ in every period he consults the expert. This cost could e.g. represent exogenously specified wages, opportunity costs of the decision maker's time spent with the expert, resources required to provide the expert access to information, or the opportunities available to the expert if he does not take up employment with the decision maker.

The assumption about the constant cost reflects circumstances in which the decision maker is a monopolist. This assumption is applicable, for example, to Presidential advisors, or to other high-profile government officers. Other applications include environments in which the salaries of advisors, experts, or consultants are fixed exogenously by law, union regulations, industry contracts, or other such institutions. Finally, the constant cost assumption is applicable if the performance of the expert is not observable to the market or if the expert has a limited time horizon and his market reputation trails his private reputation with the decision maker.

This employment cost explicitly generates a conflict of preferences between the decision maker and an incompetent expert, since the latter is of no value to the former. Consequently, the restriction formulated in Assumption 2.2, requiring the decision maker to dismiss the expert if the expected value of his future advice is 0, is redundant now because the decision maker would now strictly prefer to fire such an expert.

If a competent expert never makes mistakes, the first-best decision rule continues to be the same as in the model without cost: It calls for either never employing the expert or continuing to buy the signal as long as it has been correct and to stop doing so after a first incorrect forecast. The decision maker realizes the payoff of 0 if she never employs the expert and, otherwise, the payoff of

$$\begin{aligned} v^{FB}(N) &= \alpha\left[(N-K)p - Nc\right] \\ &\quad + (1-\alpha)\left\{\left(\tfrac{1}{2}\right)^K\left[1+\tfrac{1}{2}+\cdots+\left(\tfrac{1}{2}\right)^{N-K-1}\right]\left(\tfrac{1}{2}-(1-p)\right) - \left[1+\tfrac{1}{2}+\cdots+\left(\tfrac{1}{2}\right)^{N-1}\right]c\right\}, \end{aligned}$$

where K is defined as in the proof of Proposition 3.2 as the period after which the decision maker follows the expert's advice, i.e., K is the lowest value of N that satisfies (2) with $k = 0$.[9] Let $N^{FB}(c)$ be the lowest value of N such that $v^{FB}(N) > 0$. Note that v^{FB} satisfies the following single-crossing condition: If $v^{FB}(N) > 0$, then $v^{FB}(N') > 0$ for all $N' > N$.

Now, we observe that in this modified environment it is still true that reputational concerns vanish if $q = 1$ and N is sufficiently large.

Proposition B.1 (Vanishing Career Concerns) *Assume that the competent expert never makes mistakes. For any given p and α, there exists an integer $\check{N}_0$ such that the first-best decision rule is incentive compatible if $N \geq \check{N}_0$.*

PROOF: As the cost c does not enter the incentive constraints of the expert, the proof of incentive compatibility is identical to the proof of Proposition 3.2. The only new detail is that if $N < N^{FB}(c)$, the first best calls for no employment, which is trivially incentive compatible. ∎

[9] If $K \geq N$, $v^{FB}(N) = 0$.

Market wage

We now consider a different version of the model in which bargaining power is allocated to the expert and, at the beginning of each period, the decision maker pays him his expected value added in the given period. In addition, we assume that the expert can realize an exogenous outside option with a payoff of w_0 per period if he quits the industry. Thus, the timeline in each period in this version of the model is as follows: The state is realized. The expert is made a wage offer w_t. If he rejects it, he gets his outside option payoff of w_0 and the period ends. Otherwise, he is paid w_t and learns the value of the signal. The rest is the same as in the original model.

The assumption of a market wage being paid is in line with the literature on career concerns that considers situations in which the performance of the employee is publicly observable to a competitive market. In this market, there are no long-term contracts and the employee is free to leave the employer if there is an attractive outside offer.

In this model, the expert's salary is conditional on his performance. It is equal to 0 in periods in which the decision maker does not follow his advice; he is paid the expected increase in the decision maker's payoff otherwise. Moreover, the decision maker will obtain a payoff of 0 in each period, and so is indifferent about her behavior. The first-best decision rule, therefore, maximizes the expert's payoff.

Similarly to our previous specifications, the first-best decision rule might not be incentive compatible. To see this, let us revisit the two-period example with $q = 1$ that we use for the original model. There, the expert is valuable to the decision maker only in the second period and only if his report is correct in the first period.

Let w_2 be the value added by the expert's advice in the second period if his report is correct in the first period. To make the problem interesting, further assume that $w_2 \geq \frac{3+\alpha}{1+\alpha} w_0$. This assumption ensures that the expert's expected payoff in the first-best decision rule from truthful reporting exceeds the payoff from quitting the industry in the first period.

Now, imagine the decision maker believes that the expert reports his signal truthfully. Then, the expert's continuation payoff is again maximized by reporting 0 in the first period regardless of the signal he observes. To see this, note that the expert's wage in the first period is 0. In the second period, he obtains w_2 if his first-period report is correct and quits the industry otherwise. If he reports his signal truthfully, the probability that he is correct in the first period is $\alpha + (1 - \alpha)/2$, generating the payoff of $u^* = \frac{1+\alpha}{2} w_2 + \frac{1-\alpha}{2} w_0$. By contrast, if the expert reports 0 in the first period, his expected payoff is $u' = (1 - p)w_2 + pw_0$. By Assumption 2.1, $\frac{1+\alpha}{2} < (1 - p)$, implying $u' > u^*$.[10]

[10] Could the incentive problem be resolved by richer contracts that, for example, condition the wage on the export's report? The answer to this question depends to what extent the expert is held liable. If the expert is protected by limited liability, which we think is a natural assumption in many applications, richer contracts

The following proposition is the counterpart of Proposition 3.2.

Proposition B.2 (Vanishing Career Concerns) *Suppose that $q = 1$ and $w_0 < p$. Then, there exists $N_0(w_0)$ such that truthful reporting is an equilibrium if $N \geq N_0(w_0)$.*

PROOF: Suppose the decision maker expects the expert to report his signal truthfully. Then, the decision maker believes an expert to be incompetent if he makes an incorrect report. After that, the best response for the decision maker is to offer wages of 0, inducing the expert to reject the offer and to collect the outside option w_0. Consider now an expert whose reports have been correct in all periods preceding period t. Let α_t be the public belief at the beginning of period t. The expert's market value is given by $w(\alpha_t) = 0$ if $t \leq K$ and by $w(\alpha_t) = \alpha_t + (1-\alpha_t)\frac{1}{2} - (1-p) > 0$ if $t > K$, with K being defined as in the previous subsection. Furthermore, $w(\alpha_t)$ is increasing in t for $t > K$ and converging to p as $t \to \infty$.
Define

$$\begin{aligned} h(t,\tau) &= p\left[\alpha_t w(\alpha_\tau) + \frac{1-\alpha_t}{2}\left(\frac{1}{2^{\tau-t-1}}w(\alpha_\tau) + \left(1 - \frac{1}{2^{\tau-t-1}}\right)w_0\right)\right] + (1-p)\frac{1-\alpha_t}{2}w_0 \\ &- (1-p)\frac{1-\alpha_t}{2}\left[(1-p)^{\tau-t-1}w(\alpha_\tau) + \left(1-(1-p)^{\tau-t-1}\right)w_0\right] - p\left(\alpha_t + \frac{1-\alpha_t}{2}\right)w_0 \end{aligned}$$

for all $t \in \{1, \cdots, N-1\}$, $\tau \in \{t+1, t+2, \cdots\}$. We have that

$$h(t,\tau) = \left[\alpha_t p - \frac{1-\alpha_t}{2}\left((1-p)^{\tau-t} - \frac{p}{2^{\tau-t-1}}\right)\right](w(\alpha_\tau) - w_0).$$

In period t, the expert is willing to reveal a signal indicating the less likely state 1 if:

$$\sum_{\tau=t+1}^{N} h(t,\tau) \geq 0.$$

Observe that $\lim_{\tau\to\infty} h(t,\tau) = p\alpha_t(p - w_0)$. Hence, for each t, there exists τ' such that $h(t,\tau) \geq \epsilon$ for some $\epsilon > 0$ for all $\tau > \tau'$. Therefore, for every t there exists an $N'(t) < \infty$ such that the incentive constraint is satisfied if $N \geq N'(t)$. Let $N(t)$ be the smallest such integer.

We now define the auxiliary function $\tilde{h}(t) := h(t, t+1)$. Clearly, $\lim_{t\to\infty} \tilde{h}(t) = p(p - w_0) > 0$. Hence, there exists an integer k such that $\tilde{h}(k) > 0$. Since $\lim_{\tau\to\infty} w(\alpha_\tau) = p > w_0$, there moreover exists an integer $\hat{k}$ such that $\tilde{h}(\hat{k}) > 0$ and $w(\alpha_{\hat{k}}) > w_0$. Let $\tilde{k}$ be the smallest such integer. Now, let

trivially have no additional value: The incentive problem arises when the expert is pessimistic about the quality of his advice, but this is precisely when the value added of the expert's advice is 0. Thus, if the expert's liability is limited and hence his wages cannot be negative, paying him his market value implies that the contract is flat and pays 0 regardless of his report. Furthermore, the decision maker will not be willing to offer any contract that pays a positive wage because she must pay the expert's market value in all future periods and won't be able to recover the loss on this contract.

$N_0(w_0) := \max\{N(1), N(2), \cdots, N(\tilde{k})\}$ be the smallest value of N for which the incentive constraint is satisfied for all $t < \tilde{k}$. Then, the constraint is also satisfied for all $t \geq \tilde{k}$ and $N \geq N_0(w_0)$ because $\tilde{h}$ and $h(t, \cdot)$ are increasing, and hence $h(t, \tau) > 0$ for all $t \geq \tilde{k}$ and $\tau > t$.

Finally, it is straightforward to verify that the incentive constraint conditional on a signal indicating the more likely state 0 holds for any t if $N \geq N_0(w_0)$.

References

Benabou, R., and G. Laroque (1992): "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility," *The Quarterly Journal of Economics*, 107(3), 921–58.

Bergemann, D., and U. Hege (1998): "Venture capital financing, moral hazard, and learning," *Journal of Banking & Finance*, 22(6-8), 703–735.

——— (2005): "The Financing of Innovation: Learning and Stopping," *RAND Journal of Economics*, 36(4), 719–752.

Bourjade, S., and B. Jullien (2004): "Expertise and Bias in Decision Making," *mimeo.*

Crawford, V. P., and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50(6), 1431–51.

Dasgupta, A., and A. Prat (2006): "Financial equilibrium with career concerns," *Theoretical Economics*, 1(1), 67–93.

——— (2008): "Information aggregation in financial markets with career concerns," *Journal of Economic Theory*, 143(1), 83–113.

Effinger, M. R., and M. K. Polborn (2001): "Herding and anti-herding: A model of reputational differentiation," *European Economic Review*, 45(3), 385–403.

Ely, J. C., and J. Välimäki (2003): "Bad Reputation," *The Quarterly Journal of Economics*, 118(3), 785–814.

Foster, D., and R. Vohra (1998): "Asymptotic Calibration," *Biometrika*, 85, 379–390.

——— (1999): "Regret in the On-Line Decision Problem," *Games and Economic Behavior*, 29(1-2), 7–35.

Holmström, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169–82.

Hörner, J., and L. Samuelson (2009): "Incentives for Experimenting Agents," Mimeo.

Levy, G. (2004): "Anti-herding and strategic consultation," *European Economic Review*, 48(3), 503–525.

——— (2007): "Decision Making in Committees: Transparency, Reputation, and Voting Rules," *American Economic Review*, 97(1), 150–168.

Morgan, J., and P. C. Stocken (2003): "An Analysis of Stock Recommendations," *RAND Journal of Economics*, 34(1), 183–203.

Morris, S. (2001): "Political Correctness," *Journal of Political Economy*, 109(2), 231–265.

Olszewski, W., and M. Peski (forthcoming): "The Principal-Agent Approach to Testing Experts," *American Economic Journal: Microeconomics*.

Olszewski, W., and A. Sandroni (2008): "Manipulability of Future-Independent Tests," *Econometrica*, 76(6), 1437–1466.

Ottaviani, M., and P. N. Sørensen (2006a): "Professional advice," *Journal of Economic Theory*, 126(1), 120–142.

——— (2006b): "Reputational Cheap Talk," *RAND Journal of Economics*, 37(1), 155–175.

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

Prendergast, C., and L. Stole (1996): "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, 104(6), 1105–34.

Scharfstein, D. S., and J. C. Stein (1990): "Herd Behavior and Investment," *American Economic Review*, 80(3), 465–79.

Shmaya, E. (2008): "Many inspections are manipulable," *Theoretical Economics*, 3(3), 367–382.

Sobel, J. (1985): "A Theory of Credibility," *Review of Economic Studies*, 52(4), 557–73.

Suurmond, G., O. H. Swank, and B. Visser (2004): "On the bad reputation of reputational concerns," *Journal of Public Economics*, 88(12), 2817–2838.

Trueman, B. (1994): "Analyst Forecasts and Herding Behavior," *Review of Financial Studies*, 7(1), 97–124.

Yüceer, U. (2002): "Discrete convexity: convexity for functions defined on discrete spaces," *Discrete Applied Mathematics*, 119, 297–304.