Little White Lies–The Value of Inconsequential Chatter*

Nicolas Klein† Tymofiy Mylovanov‡

This version: April 15, 2011

Abstract

This paper deals with the problem of providing adequate incentives to an expert who might be tempted to conceal his true opinion because of his desire to appear competent. We show that if a competent expert never makes mistakes, the incentive problem will disappear if the interaction lasts long enough. However, if a competent expert occasionally makes mistakes, the opposite obtains: There will always arise an incentive problem if the time horizon is sufficiently long. We furthermore demonstrate that the decision maker can address the incentive problem by letting the expert accumulate some private information about his ability, and that doing so is optimal if the competent expert does not make mistakes too often.

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

JEL Classification Numbers: C73, D83.

*We are grateful to Dirk Bergemann, Andreas Blume, Larry Blume, Matthias Fahn, Johannes Hörner, Stephan Lauermann, Qingmin Liu, Lucas Maestri, George Mailath, Benny Moldovanu, Stephen Morris, Mallesh Pai, Andrew Postlewaite, Andrea Prat, Sven Rady, Phil Reny, Luca Riggotti, Larry Samuelson, Satoru Takahashi, and Rakesh Vohra for useful comments and discussions. 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, Penn State University; email: txm41@psu.edu.

> THE choice of servants is of no little importance to a prince, and they are good or not according to the discrimination of the prince. And the first opinion which one forms of a prince, and of his understanding, is by observing the men he has around him; and when they are capable and faithful he may always be considered wise, because he has known how to recognize the capable and to keep them faithful.
>
> Nicolo Machiavelli, “The Prince,”
> Ch. XXII “Concerning The Secretaries of Princes”

1 Introduction

This paper deals with a decision maker (she) and an expert motivated by career concerns (he), who interact repeatedly over a finite time horizon. For a specific motivation consider, for instance, a legislator who can consult a pollster to find out whether the people she represents would prefer a vote in favor of, or against, a particular bill under consideration. Not fully certain about the quality of his research, the pollster might decide to play it safe and distort his report toward the commonly known ideological bent of the legislator's district. The legislator's goal meanwhile is twofold: She wants to make the best possible use of the advice she gets in the current period, while at the same time learning about her expert's competence, as this will help her make better decisions in the future.

More broadly, our analysis applies to many instances when a decider who faces a sequence of multiple decision problems can elect to seek the help of outside advisors whose quality is initially unknown: A firm trying to gauge the demand for a new product may look toward outside consultants for advice. People filing their income tax returns may hire a tax consultant. Or a college student having set aside an afternoon to study for a test may decide to spend the time with a tutor.

We focus on a situation in which initially both parties are equally uncertain about the expert's competence;[1] at the end of each period, though, both parties will publicly observe if the expert's prediction will have come to pass, and they will update their respective opinions about his competence accordingly. Crucially, though, the decision maker only observes the expert's report; she does not observe the expert's information that led him to make the report.

If the expert 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 expert *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.

[1] Focussing on the case of a common initial prior allows us to isolate the incentive problems solely caused by the expert's career concerns. This assumption standard in the literature (see e.g. Ottaviani and Sørensen (2006a, 2006b)).

The observation that an expert's reputational concerns can have adverse effects on his incentives to be truthful has been made by Morris (2001).[2] In his paper, the expert's reputation pertains to his preferences rather than his competence. The important message of his work is that strong reputational concerns make meaningful communication impossible. We show that this conclusion is reversed in our environment, where the strength of the expert's reputational concerns can be measured by the potential length of the relationship: If the competent expert never makes mistakes, incentives to report truthfully are restored as the number of periods grows sufficiently large (Proposition 3.2). Thus, if the time horizon is long enough, forward-looking reputational concerns will discipline the expert's behavior to the point of completely counterbalancing the harmful myopic ones.

Surprisingly, though, if the competent expert occasionally makes mistakes, this result 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 limited in such a way that they may always be overcome by harmful myopic ones if the time horizon is sufficiently long.

The fact that effective communication can be impeded by an expert's concern for appearing competent has been observed in the context of *single-decision* environments, e.g. by Trueman (1994), and Ottaviani and Sørensen (2006a, 2006b). Prat (2005) shows that a similar effect arises if the expert is delegated the decision rights over the actions.[3]

Our paper takes a new angle in that we consider a *multi-decision environment* in which the expert's private belief about his competence may over time diverge from that of the decision maker. This gives rise to the *bad news effect*: If the expert distorts his information 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 expert never makes mistakes, incentives for truthful reporting are restored as the number of periods grows sufficiently large (Proposition 3.2). Formally, the difference between the continuation payoff after the expert 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 expert's continuation payoff after a mistake is zero as it always results in termination.

If a competent expert occasionally makes mistakes, by contrast, the forward-looking reputational effect is bounded. The reason for this is that the expert's payoff from truth-telling now converges as the number of periods increases. Yet, at the same time, the strength of the myopic reputational effect is increasing in the time horizon, because the decision maker is willing to

[2]See Morgan and Stocken (2003) for an alternative model with uncertainty about the expert's preferences. The early model of preference-based repuational concerns in communication is Sobel (1985).

[3]Reputational concerns about agents' competence are also studied in Scharfstein and Stein (1990), Prendergast and Stole (1996), and in Suurmond, Swank and Visser (2004).

tolerate a larger number of mistakes as the time horizon increases. Thus, the bad news effect is not strong enough to restore truth-telling incentives when the expert has grown sufficiently pessimistic about his ability.

The introductory quote by Machiavelli would suggest that the decision maker should concern herself with distinguishing between good experts and bad experts,[4] and hence with creating incentives for the expert to be forthright with his information. More generally, the implementation of truthful information transmission has been one of the main concerns of the literature on cheap-talk communication (e.g., Crawford and Sobel (1982), Battaglini (2002), Levy and Razin (2007), Ambrus and Takahashi (2008), Ambrus and Lu (2010), Esö and Fong (2010), Ivanov (2011)). We show that even if the bad news effect is not sufficiently strong to restore incentives, the decision maker can still induce the expert to report his information truthfully by threatening to punish him for forecasting the *a priori* likelier outcome even if this forecast should turn out to be correct (Proposition 4.1).

However, the decision maker cares about truth-telling only insofar as it helps her make better decisions; indeed always inducing the expert to tell the truth need not be her best option. For the case in which a good expert never makes mistakes, we construct the optimal equilibrium and show that whenever the bad news effect is not strong enough to obviate all incentive problems, these are best addressed by letting the expert gain some private knowledge about his abilities in the first few periods of interaction (Proposition 4.3).[5] Thereafter, the expert 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 expert, 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 expert's occasional white lies avoids the decision maker the cost of sometimes losing the valuable services of an expert whose only fault has been correctly to predict the expected. Our analysis would thus suggest that, in certain circumstances, some amount of tolerance of inconsequential chatter and the occasional white lie might outperform a single-minded focus on truthful information transmission.

The equilibria we construct remain equilibria even if a competent expert occasionally makes mistakes. However, the problem becomes significantly more complex in this case, which makes it difficult to compare the performance of different decision rules.

Negative reputational effects also appear in Ely and Välimäki's environment (2003), in

[4]In a sense, our decision maker's problem is akin to that of an economic agent operating a two-armed bandit machine with one safe arm (policy based on the prior beliefs) and one risky arm (consulting an expert), whose expected payoff is initially unknown but can potentially be learnt through use over time. See Bergemann and Välimäki (2008) for an overview of this literature.

[5]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.

which the expert takes an action rather than offering advice. In their model, the expert is informed about his type and the preferences of the good expert, who is perfectly informed about the state of the world in each period, are completely aligned with those of the principal; meanwhile, the bad expert is incompetent, and has myopic incentives to choose a certain action in each period. They show that, as in Morris (2001), the value of the expert vanishes as parties become sufficiently forward-looking, provided the agent faces a sequence of short-lived principals. Otherwise, the principal can achieve his full-information payoff.

Our paper is also related to the literature on expert testing.[6] The paper most closely connected to ours in this literature is Olszewski and Peski's (forthcoming) infinite horizon principal-agent model. In their model, experts privately know their type and a competent expert knows the probability distribution of future states. The key difference is that in our model the expert privately observes some information about the realized state, and hence about his type, in each period; in Olszewski and Peski (forthcoming), by contrast, the expert's type describes his knowledge about the distribution of states, and no additional private information is accumulated over time. Olszewski and Peski find that the first best can be approximated if the parties become infinitely patient.

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

Finally, there is a number of less directly related papers: The idea that agents may want to conform to what the principal thinks should likely be their information also appears in Prendergast (1993), where the quality of the expert's information depends on his effort choice. Dasgupta and Prat (2006, 2008) show that financial traders' career concerns relating to their reputation for competence can increase trading volumes and prevent asset prices from reflecting fundamental values. Levy (2007) studies career concerns in committee decision making. Gerardi and Maestri (2009) study a cheap-talk model without career concerns, in which the agent has to be incentivized to exert costly effort in order privately to observe a sequence of signals about the decision-relevant state of nature.

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 decision rules, whereas the second-best decision rules are studied in Section 4; Section 5 concludes. The proofs omitted in the main text are provided in the appendix.

[6]See e.g., Foster and Vohra (1998, 1999); the most general results are obtained in Olszewski and Sandroni (2008) and Shmaya (2008).

2 Model Setup

We study the simplest model that formalizes the expert'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 expert before making a policy choice. The expert 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 expert 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. 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 expert's *competence*.

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

The core issue that our model is built to address is the following: If the expert were non-strategic, the decision maker would like to employ the expert as long as she believed he could still be of help, possibly further down the road. Yet, if this decision rule were in fact followed, the expert might have incentives to suppress a priori unlikely information to maximize his chances of appearing competent; this would slow down learning about the expert's quality and thus render his advice less valuable. Hence, to rule out uninteresting cases, we impose the following

Assumption 2.1 *It is commonly believed that* $\beta_1 := \alpha q + (1 - \alpha)/2 < 1 - p$*;*

i.e. the decision maker obtains a higher payoff if she follows her prior beliefs than if she follows the signals of an expert with reputation α. Simultaneously, this assumption implies that an expert 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 expert. If he is employed, the expert 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 expert'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 expert 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 expert 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 expert. This cost could e.g. represent exogenously specified wages, opportunity costs of the decision maker's time spent with the expert, or resources required to provide the expert 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 expert's career concerns would have no impact in our model because it would be optimal for the decision maker simply never to fire the expert.

Thus, 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 expert'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 expert 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 expert 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.

$$\beta_t > 1 - p. \tag{1}$$

If the expert's report is not sufficiently informative or if the expert 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 First Best

As our first-best benchmark, we consider a hypothetical environment in which the expert's signals are observed by the decision maker.[7] Let $\alpha_N(k)$ denote the posterior belief that the

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

expert 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 expert'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

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

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 expert's signal is valuable for the decision maker in the last period.

Definition The first-best decision rule

1. employs the expert until his reports have disagreed with the state $\kappa + 1$ times;
2. implements a policy equal to the expert's report if $\alpha_t > \frac{1-2p}{2q-1}$ and policy 0 otherwise.

If the expert's reports are truthful, this rule is a best response for the decision maker because it maximizes her payoff and retains the expert 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 expert's career concerns. Furthermore, the decision maker's payoff if she follows the first-best decision rule and the expert 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 expert entails an 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 expert'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 expert 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 expert believes that the state $\omega_1 = 0$ is more likely. Thus, the probability of employment in the next period is maximized by reporting $\hat{s}_1 = 0$. As a result, the expert'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 expert never makes mistakes,[8] 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 expert 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 expert believes that he is more likely to be competent. Moreover, if the report is incorrect, the expert 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 expert evaluates his future payoffs conditional on different events.

Proposition 3.2 (Vanishing Career Concerns) *Assume that the competent expert never makes mistakes. For any given p and α, there exists an integer N_0 such that the first-best decision rule is incentive compatible if and only if $N \geq N_0$.*

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 N_0 such that $N(t) \leq N_0$ 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 expert is employed toward the end of the relationship under the first-best decision rule, then his reputation is necessarily high, the expert 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 expert is always correct. However, if the competent expert 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 expert 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 expert 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

[8] This assumption implies that the decision maker will learn that the expert is incompetent for sure whenever he is expected to tell the truth and his report is inconsistent with the realized state. Similar full-revelation assumptions are commonly made, e.g. in the principal-agent models of Gerardi and Maestri (2009), Bergemann and Hege (1998, 2005), Hörner and Samuelson (2009).

this rule, the expert'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 expert 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 expert will prove himself to be sufficiently competent to become valuable for the decision maker. However, after a mistake, the expert is no longer willing to report his signal truthfully. If $N = 2$, this does not matter as the expert 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 expert is of no value to the decision maker. As a matter of fact, the following Proposition shows that the first-best decision rule will always be incentive incompatible if N exceeds a certain threshold.

Proposition 3.3 (Persistent Career Concerns) *Suppose the competent expert 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 result does not address how much of a loss the decision maker will incur as a result of the incentive problem that precludes the first-best outcome. Although an incentive problem arises with positive probability after some histories, the set of these histories might become small as $N \to \infty$, and the decision maker might be able to remedy the problem by slightly modifying the decision rule. Unfortunately, there is no simple solution. Recall that in the environment in which a competent expert never makes mistakes, the positive result is due to the bad news effect. This effect implies that, in the first-best decision rule, the expert's expected continuation payoff from telling the truth diverged in N, while his payoff from lying stayed bounded. This is true for any belief about the expert's competence. However, this is no longer true when a competent expert may also make mistakes. Now, the continuation payoff from telling the truth at history nodes at which the expert is fired after one (or say k) additional mistake(s) does not diverge in N. As a result, there exists an $\alpha^* > 0$ such that truth-telling is not incentive compatible for any N if the expert's belief about his competence is lower than α^*. The difficulty is that the probability that a belief less than α^* is reached under truth-telling is fully determined by the prior belief; this probability does not vanish as $N \to \infty$. Furthermore, as $N \to \infty$, an expert with a reputation lower than α^* might still turn out to be valuable to the decision maker in the distant future, so that she should continue to employ him in the first-best rule.

This concludes our section on the first-best decision rule. We study the second-best decision rules in the next section.

4 Second Best

4.1 Fully Revealing Equilibria

Even if the first-best decision rule is not incentive compatible, the decision maker could still try and enforce truth-telling by threatening to fire the expert for good with some probability after he has correctly forecast the more likely state. This threat is in turn made credible by the expert's off-path threat to babble and send but uninformative messages if he continues to be employed and the decision maker does not fire him when she is supposed to do so.[9]

In order to ensure that a deviation by the decision maker from firing the expert with the probability specified in equilibrium is observable, the game can be modified to allow for jointly controlled lotteries (Aumann and Maschler 1995).[10]

We assume that after the decision maker takes an action but before the state is publicly observed, both players simultaneously send messages from the set $[0, 1]$ and these messages immediately become common knowledge. Now imagine, for instance, that the decision maker is supposed to fire the expert with a probability of one half if his report is incorrect. Then, in equilibrium, both players randomize between two messages with equal probability, e.g., 0 and 1. If the messages coincide, which happens with probability one half, the decision maker continues to employ the expert. Otherwise, the expert is fired. These strategies achieve the desired probability of firing and, moreover, a deviation by the decision maker to any strategy that employs the expert with a different probability is publicly observable. Furthermore, a deviation to a different message strategy by a single player does not affect the expert's employment probability, and hence the reporting strategies are mutually best responses. A similar construction is possible for any firing probability, although it will typically require more complex message strategies.

Definition An equilibrium is *fully revealing* if on the equilibrium path the expert reports his signals truthfully.

Of course, there are uninteresting fully revealing equilibria in which the expert is never hired, or is hired only in the last period. The following result, though, shows that non-trivial fully revealing equilibria can also be constructed. Indeed, we construct an equilibrium in which the expert is fired right after his first mistake; moreover, after he has correctly forecast the initially more likely state he is fired with some positive probability which is chosen so as to give him incentives to report truthfully. This construction is quite natural for $\kappa = 0$. For larger κ, this construction is still an equilibrium, even though more complicated fully revealing equilibria might perform better.

[9]This is incentive compatible for the expert, because the decision maker does not condition his continued employment on the reports.

[10]Jointly controlled lotteries have been used to construct equilibria in cheap talk environments without career concerns, e.g., in Aumann and Hart (2003), Krishna and Morgan (2004), Forges and Koessler (2008).

Proposition 4.1 (Fully Revealing Equilibrium) *There exists a fully revealing equilibrium in which the expert is hired in the first period and is always retained along the equilibrium path after a correct report* 1*, retained with positive probability after a correct report* 0*, and fired after an incorrect report.*

PROOF: We construct a fully revealing equilibrium with the following strategies.

Messages. The expert reports the true state unless he has observed a deviation by the decision maker, in which case he reports 0 in each period he is consulted. Let A_1 and A_2 respectively denote the messages sent by the expert and the decision maker in the jointly controlled lottery. Both players draw their messages A_1 and A_2 from the uniform distribution on $[0, 1]$.

Employment. On the equilibrium path, the decision maker fires the expert whenever the report is incorrect and retains the expert after a correct report 1. After a correct report 0, the expert is retained if and only if $A_1 - A_2 \in [0, \rho^s_{0,t}]$ or $A_2 - A_1 \in [1 - \rho^s_{0,t}, 1]$, where $\rho^s_{0,t}$ is constructed by induction in descending order of periods: In each period set its value equal to the maximal value with which the expert can be retained after correctly reporting 0 such that reporting truthfully in this and future periods is incentive compatible given the players' continuation strategies.

If the decision maker has deviated from her employment strategy, she subsequently never consults the expert.

Mutual best response property. First, consider a deviation by the decision maker from her employment strategy. This deviation is commonly known to trigger uninformative reports on the expert's part. This makes it optimal for the decision maker not to employ the expert in the future. In turn, even if the expert is consulted, reporting 0 is a best response as he expects not to be consulted ever again.

Second, neither player can affect his expected probability that the expert will be retained after a correct report of 0 by unilaterally deviating to a different distribution over his respective message A_i. Hence, these strategies are mutually best responses.

Next, by construction of our message strategies, the expert is retained after a correct report 0 with probability $\rho^s_{0,t}$. Again, by construction, the probability is chosen such that truthful reporting is a best response for the expert.

Moreover, employing an expert who has not made an incorrect report is a best response for the decision maker by Assumption 3.1. ∎

4.2 Equilibrium With Endogenous Private Information

A quite natural way for the decision maker to handle the expert'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 expert 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 proposition that follows. First, assume the decision maker follows the first-best decision policy. Now, let t^{FB} be the earliest period such that an expert who has observed and reported only correct signals will henceforth find truthful reporting optimal.[11] (Clearly, if $t^{FB} = 1$, the first-best decision rule is incentive compatible. Furthermore, $t^{FB} \leq N - 1$ because the expert is indifferent about his report in the last period.)

Proposition 4.2 (Equilibrium With A Grace Stage) *There exists an equilibrium in which no information is transmitted, and the expert is never fired during the first t^{FB} periods; thereafter, the expert 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 expert'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 expert 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.[12]

The decision maker's equilibrium strategy calls for (0) not hiring the expert in those periods τ such that there exists a period $\tilde{\tau} < \tau$ in which the expert has given an incorrect forecast and $\tilde{\tau} > t^{FB}$, or in which the expert has reported 1 and $\tilde{\tau} \leq t^{FB}$. In all other periods, she employs the expert.

These strategies are mutually best responses by the definition 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: Suppose the expert has privately learned that he is of the bad type; he then maximizes his expected employment duration by always reporting state 0. A decision maker who knew that she could not rely on the expert's advice would also optimally stick with her prior and implement policy 0. By contrast, in any fully revealing equilibrium, a good expert will be fired with positive probability, leading to worse policy decisions in expectation from an ex-ante point of view. Thus, at the optimum,

[11] That is, the expert'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, independently of the history of signals in periods $t > t^{FB}$.

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

the first-best quality of policy decisions is achieved thanks to a longer ex-ante expected duration of employment than in the first best. This is in contrast e.g. to Gerardi and Maestri (2009), where employment of an expert is costly, and hence there is a downward distortion of his expected employment duration in the optimal contract.

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 expert 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 expert 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 expert, private information would rather tend to hurt the principal.[13]

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. Still, it violates condition 1. of our definition of the first best, as the expert 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 expert entails a small cost for the decision maker). Of course, if the decision maker incurred such a (small) cost for employing the expert, she would prefer firing a bad expert as quickly as possible. As it turns out, it is impossible to achieve v^{FB} while employing the expert 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 4.3 (Second-Best Optimum) *If $\kappa = 0$, the decision maker's ex-ante expected payoff in the equilibrium identified in Proposition 4.2 is equal to v^{FB}. 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 expert'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 v^{FB} in which the expert is employed for fewer periods in expectation. In order for the principal to achieve an ex-ante expected value of v^{FB}, it must be the case that a good expert is never fired; i.e. in such an equilibrium, the expert 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 4.2, it must be the case that some information on the agent's type will be transmitted in period t^{FB} or earlier. Yet, to induce the expert to tell the truth with some positive probability in period t^{FB} or earlier, by Assumption 2.2, the decision maker has to fire

[13] 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 expert 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.

the expert with some positive probability even after he has been correct. This in turn implies that a good expert will be fired with positive probability. Hence, the decision maker makes worse policy decisions in expectation, and thus her payoff is bounded away from v^{FB}. ∎

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. For a trivial instance of the latter effect, recall our Example on page 8. There, the first best can be achieved by having a "grace period" in $t = 2$ if the expert makes a mistake in $t = 1$, which here, as an artefact of the simple three-period structure, boils down to never firing the expert after any history. We leave a rigorous exploration of these issues outside the scope of this paper.

5 Conclusion

We have investigated the dynamic interaction between a decision maker and an expert 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 expert may have an incentive strategically to manipulate the cheap-talk relay of his signal to the decision maker. We have shown that if a competent expert never makes mistakes and the number of periods is large enough, the expert's career concerns vanish, and the first best becomes implementable; however, the opposite is true if a competent expert occasionally makes mistakes. Moreover, we have shown that the decision maker can address the incentive problem by letting the expert accumulate some private information about his ability; doing so is optimal if a competent expert never makes mistakes.

In our model, the decision maker can only set incentives by either retaining or firing the expert. 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 expert. We would conjecture that our decision maker would want to do so if she was faced with an optimistic expert,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 expert's learning about his type, she would not be able completely to shut it down, as the expert 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 experts, so as to expedite their learning process. We leave a full exploration of these issues for future work.

Appendix

Proof of Proposition 3.2

Suppose the decision maker pursues the first-best policy of immediately firing the expert if, and only if, the expert 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:

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

To understand the right-hand side of the incentive constraint, the reader should note that if, upon lying, the expert 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 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 $N_0 \geq 2$.

It is left to check that 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, this constraint also holds for all $N \geq N_0$. ∎

Proof of Proposition 3.3

Fix arbitrary parameters α, p and $q < 1$. Let h^* be a history such that (1) the expert 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 expert's belief about his competence, and $K = N - t$ the remaining number of periods at h^*. Then, this condition can be expressed as

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

or, equivalently,

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

The LHS is increasing in K and converges to $\alpha'(2q-1)\left(\frac{p}{1-q}+\frac{1-p}{2p}\right)$ from below, while the RHS 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 expert'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:

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

From (A.4), we have that as N diverges, the RHS 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 RHS diverges in N and hence $N-\kappa$ diverges. ∎

References

Ambrus, Atilla and Shih-En Lu, "Robust fully revealing equilibria in multi-sender cheap talk," Mimeo 2010.

Ambrus, Attila and Satoru Takahashi, "The multi-sender cheap talk with restricted state spaces," *Theoretical Economics*, 2008, *3*, 1–27.

Aumann, Robert J. and Michael Maschler, *Repeated Games with Incomplete Information*, MIT Press, 1995.

——— **and Sergiu Hart**, "Long Cheap Talk," *Econometrica*, November 2003, *71* (6), 1619–1660.

Battaglini, Marco, "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica*, 2002, *70*, 1379–1401.

Bergemann, Dirk and Juuso Välimäki, "Bandit Problems," in Steven Durlauf and Larry Blume, eds., *The New Palgrave Dictionary of Economics*, 2 ed., McMillan Press, 2008.

—— **and Ulrich Hege**, "Venture capital financing, moral hazard, and learning," *Journal of Banking & Finance*, August 1998, *22* (6-8), 703–735.

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

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

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

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

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

Esö, Peter and Yuk-fai Fong, "Wait and See: A Theory of Communication over Time," Mimeo 2010.

Forges, Françoise and Frédéric Koessler, "Long persuasion games," *Journal of Economic Theory*, November 2008, *143* (1), 1–35.

Foster, Dean and Rakesh Vohra, "Asymptotic Calibration," *Biometrika*, 1998, *85*, 379–390.

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

Gerardi, Dino and Lucas Maestri, "A Principal-Agent Model of Sequential Testing," Mimeo 2009.

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

Hörner, Johannes and Larry Samuelson, "Incentives for Experimenting Agents," Mimeo 2009.

Ivanov, Maxim, "Dynamic Informational Control," Mimeo 2011.

Krishna, Vijay and John Morgan, "The Art of Conversation: Eliciting Information from Experts Through Multi-Stage Communication," *Journal of Economic Theory*, 2004, *117* (2), 147–79.

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

—— **and Ronny Razin**, "On the limits of communication in multidimensional cheap talk: a comment," *Econometrica*, 2007, *75*, 885–893.

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

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

Olszewski, Wojciech and Alvaro Sandroni, "Manipulability of Future-Independent Tests," *Econometrica*, November 2008, *76* (6), 1437–1466.

___ **and Marcin Peski**, "The Principal-Agent Approach to Testing Experts," *American Economic Journal: Microeconomics*, forthcoming.

Ottaviani, Marco and Peter Norman Sørensen, "Professional advice," *Journal of Economic Theory*, January 2006, *126* (1), 120–142.

___ **and** ___ , "Reputational Cheap Talk," *RAND Journal of Economics*, Spring 2006, *37* (1), 155–175.

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

Prendergast, Canice, "A Theory of "Yes Men," *American Economic Review*, September 1993, *83* (4), 757–70.

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

Scharfstein, David S and Jeremy C Stein, "Herd Behavior and Investment," *American Economic Review*, June 1990, *80* (3), 465–79.

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

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

Suurmond, Guido, Otto H. Swank, and Bauke Visser, "On the bad reputation of reputational concerns," *Journal of Public Economics*, December 2004, *88* (12), 2817–2838.

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