

Reputational cheap talk

Marco Ottaviani*

and

Peter Norman Sørensen**

We analyze information reporting by a privately informed expert concerned about being perceived to have accurate information. When the expert's reputation is updated on the basis of the report as well as the realized state, the expert typically does not wish to truthfully reveal the signal observed. The incentives to deviate from truthtelling are characterized and shown to depend on the information structure. In equilibrium, experts can credibly communicate only part of their information. Our results also hold when experts have private information about their own accuracy and care about their reputation relative to others.

1. Introduction

■ To foster their careers, experts typically seek favorable evaluations. For example, consider an analyst consulted by an investor about the value of a stock. The investor evaluates the accuracy of the analyst's information by cross-checking the analyst's report with the realized profitability of the investment. Analysts who are believed to have access to more accurate information are rewarded by the market. Similarly, managers and business consultants with better reputations can charge more for their services. Company directors and board members who are believed to be effective are more likely to be reappointed or hired by other companies. Politicians considered to be better informed are more likely to be re-elected and so derive increased private benefits. The model of reputational cheap talk developed in this article can be applied to these and other situations in which individuals are concerned about their reputation for being well informed.

Our model features an expert who observes a private signal (s) about a state of the world (x). The amount of information contained in this signal depends on the expert's ability type (t). After observing the signal (and possibly the ability), the expert reports a message (m) to an evaluator. The evaluator later observes the state and combines it with the message to update the belief regarding the expert's ability. This belief (or reputation) determines the expert's payoff, which the expert aims to maximize. This is a cheap-talk (or costless-signalling) game, in which the expert

* London Business School; mottaviani@london.edu.

** University of Copenhagen; peter.sorensen@econ.ku.dk.

Some of the results presented in this article were previously contained in UCL Discussion Paper no. 99–05 (“Notes on Reputational Cheap Talk”) and Copenhagen Economics Discussion Paper no. 02–05 (“Professional Advice: The Theory of Reputational Cheap Talk”). These earlier papers also contained the detailed analysis of an example of this model with a multiplicative linear experiment (Example 1c in this article). That material and the analysis of dynamic extensions of the model now appear separately in the companion article “Professional Advice.” We thank the participants of ESSET 2004 (Gerzensee) and especially the Editor and referee for helpful comments.

(sender) does not bear a direct cost from the message sent, but cares about the induced response of the evaluator (receiver).^{1,2}

In the context of this model, we aim at addressing the following questions:

- (i) When is the expert's concern for ability compatible with truthful reporting?
- (ii) In which direction does the expert wish to bias the report, when believed to be truthful?
- (iii) What reporting strategy should the expert be rationally expected to use in equilibrium?
- (iv) What is the effect of the expert's knowledge about own ability?
- (v) What happens in the presence of competition among experts?

The main result of this article is that with a generic continuous information structure, the equilibrium cannot be fully revealing. Contrary to naive intuition, experts wishing to be perceived as accurate do not truthfully reveal their private information. In any generic model in which the ability parameterizes the signal's informativeness about the state, the information about the state is necessarily intertwined with that about ability. There are then two incentives at play. On the one hand, the dependence of the signal on ability gives an incentive to manipulate the signal's report, but on the other hand this incentive is mitigated by the signal's dependence on the state.

To understand why truthtelling is typically incompatible with equilibrium, suppose that the evaluator conjectures that the expert truthfully reveals the signal. Note that the ability determines the relation between the signal and the state of the world. Then, some signal-state observations made by the evaluator must result in favorable posterior reputation about ability. In a putative fully separating equilibrium, the signal and the realized state are informative about ability. This gives the expert an incentive to manipulate the report of the signal in order to generate a better reputation. The expert will then want to pretend to possess the signal that gives rise to the highest expected reputational payoff, and this is typically different from the signal actually observed. The expert's incentive to lie destroys any fully separating equilibrium in which the expert's signal can be inferred from the message reported. We prove that only degenerate specifications of the information structure are so balanced that the true signal can be honestly reported in equilibrium (Proposition 1).

This finding is reminiscent of Crawford and Sobel's (1982) result in the canonical model of partisan advice, but it is driven by different forces. In their setting, the impossibility of full revelation follows immediately from the presence of conflict of interest between the sender and the receiver. By contrast, in our setting this impossibility crucially depends on the features of the information structure.³ To illustrate this point, we showcase two somewhat degenerate information structures that permit truthtelling (Propositions 2 and 3). In order to better understand the incentive to deviate from truthtelling, we also provide an analytical characterization of its direction (Proposition 4). In our focal examples the deviation incentive leads experts to bias their reports toward the *a priori* expected (Proposition 5). This is because experts fear that extreme predictions are too likely to be seen as the product of noisy, uninformative signals.

More generally, in equilibrium the message space is endogenously coarse. Experts who desire to impress a rational audience are then unable to communicate all the information they have. When experts with different information rank differently the evaluation of ability following the various messages sent, it is possible for some information to be communicated in equilibrium.⁴

¹ In the cheap-talk model of Crawford and Sobel (1982), the receiver cannot commit to take any decision other than the *ex post* optimal one given the information communicated by the sender. In contrast, in the case of delegation (Holmström, 1984), the receiver is able to commit *ex ante* to a decision rule.

² See Sobel (1985), Benabou and Laroque (1992), and Morris (2001) for models of reputation building about preferences, rather than quality of information possessed. While those articles study dynamic games, we look at the stage cheap talk game with exogenous reputational concerns.

³ There are two key differences between partisan and professional advice. First, sender and receiver have different objectives in the two models. Second, our receiver has access to an additional source of information (the realized state) before taking the action (evaluation of the sender).

⁴ Similarly, Crawford and Sobel (1982) find that some communication is possible when the sender and the decision maker have sufficiently congruent preferences.

In any case, a reported message pools many signals and is therefore less precise than the sender's true signal.

Intuitively, the incentive to deviate from truthful revelation should depend on whether the sender knows her own ability. The incentive to deviate toward what the receiver expects to hear should be diminished when ability is privately known, because higher-ability experts tend to have more extreme beliefs. We clarify through some examples that this intuition depends in a subtle way on whether the sender is expected to truthfully report the signal or the posterior belief. If the evaluator expects truthful reporting of the signal, the expert is still conservative, although less so the more able she is. If instead the evaluator expects truthful reporting of the posterior mean of the expert's posterior belief, experts wish to exaggerate their beliefs (Proposition 6). Contrary to what happens when the expert does not know her own ability, with known ability there exists always an informative equilibrium in our examples (Proposition 7).

We also investigate the effect of competition among experts on their reporting incentives. In a winner-take-all specification, experts might well have incentives to differentiate their messages. Instead, we show that if the experts continue to have von Neumann-Morgenstern payoffs and conditionally independent signals, a multi-expert setting is the direct sum of equivalent single-expert settings (Proposition 8).

In the Bayesian statistics literature, Bayarri and DeGroot (1988) and (1989) were the first to analyze the incentive to manipulate information reports in order to gain influence. They posited that the weight given to an expert is proportional to an expert's prior weight and the predictive density that the expert had assigned to the realized outcome.⁵ In their setting, experts maximize their own weight by reporting a predictive distribution different from their posterior distribution. We depart from Bayarri and DeGroot in two important ways. First, rather than assuming an ad hoc updating rule for the weights, we follow the lead of Holmström (1999) by positing that the evaluator makes optimal use of all information available to form the posterior belief on the informativeness of the expert's signal.⁶ Second, we not only characterize the incentives to deviate from honest reporting, we also study the equilibrium of the game.

Our work builds more directly on the single-agent model of Scharfstein and Stein (1990).⁷ In their model signals, states, and ability types are binary and the signals of better informed experts are conditionally more correlated. They observed that the second of two agents could not communicate honestly, but they attributed this finding mainly to the assumption of differential conditional correlation. Here, we consider instead experts with conditionally independent signals and discover that honesty is impossible under very general conditions.⁸

The article is organized as follows: In Section 2 we set up the model for the baseline case of an expert who does not know her own ability. Section 3 shows the generic impossibility of truthful revelation. Section 4 investigates the deviation incentives. In Section 5 we discuss the equilibrium predictions of the model, and in Section 6 consider the case of known ability. Section 7 extends the model to the case with multiple experts who care about the perception about their relative ability. All proofs are collected in the Appendix.

2. Model

■ We consider a communication game played by an expert (sender) and an evaluator (receiver) with the following timing. First, the expert receives a signal realization that depends on state and ability. Second, the expert sends a message to the evaluator. Third, the evaluator observes the state

⁵ This happens naturally if a linear opinion pool is used (e.g., see Genest and Zidek, 1986).

⁶ Section 3 of Holmström (1999) contains the first formulation of a reputational model where more able managers have access to a more precise signal about an investment opportunity. For a general analysis of the moral hazard problem presented instead in the first part of Holmström's article, see Dewatripont, Jewitt, and Tirole (1999).

⁷ Departing from Holmström (1999), they assumed that the state of the world is eventually realized regardless of the report. We also assume that the sender's report does not affect the state or what the receiver can observe about it.

⁸ See also the discussion in Ottaviani and Sørensen (2000).

and updates the belief about the expert's ability. We now explain in more detail the information structure and the payoffs of the players.

An expert of ability type $t \in T \subseteq \mathbb{R}$ privately receives an informative signal $s \in S$ on the state of the world $x \in X$ with conditional probability density function (p.d.f.) $f(s | x, t)$. The ability parametrizes the informativeness of the expert's signal. State and ability are assumed to be statistically independent, with common nondegenerate prior beliefs $q(x)$ on state and $p(t)$ on ability. Until Section 6, we assume that the expert does not know her own ability type t . After observation of the nonprovable signal s , the expert decides which message $m \in M$ to send. As explained below in more detail, the message space is arbitrarily rich and determined only as part of an equilibrium. A strategy of the expert is a mapping from signals into messages. The conditional probability that m is sent following signal s is denoted by $\varphi(m | s)$. When representing a pure strategy, the message sent after receiving signal s is denoted by $m(s)$.

The evaluator observes the message sent by the expert as well as the eventual realization of the state x . The evaluator's job is to compute the posterior reputation of the sender $p(t | m, x)$.^{9,10} To do so, the evaluator forms a conjecture $\hat{\varphi}$ on the strategy used by the sender. Given the conjecture, the evaluator computes the chances $\hat{f}(m | x, t) = \int_S \hat{\varphi}(m | s) f(s | x, t) ds$ and $\hat{f}(m | x) = \int_T \hat{f}(m | x, t) p(t) dt$. The posterior reputation is then calculated according to Bayes' rule to be $p(t | m, x) = p(t) \hat{f}(m | x, t) / \hat{f}(m | x)$.¹¹

Regardless of the privately observed signal, the expert wishes to induce the evaluator's most favorable beliefs $p(t | m, x)$.¹² The expert's preferences over posterior reputations are represented by the strictly increasing von Neumann-Morgenstern utility function $v(t)$.¹³ The reputational payoff of message m in state x is

$$W(m | x) \equiv \int_T v(t) p(t | m, x) dt. \quad (1)$$

The preference ordering over reputation for expertise is therefore common across types. When sending the message, the sender does not know the state that the evaluator will observe and use to update the sender's reputation.¹⁴ The expected reputational payoff for a sender with signal s who sends message m is

$$V(m | s) = \int_X W(m | x) q(x | s) dx, \quad (2)$$

where the expert's posterior belief on the state x conditional on receiving signal s is given by Bayes' rule as $q(x | s) = f(s | x) q(x) / f(s)$, with $f(s | x) = \int_T f(s | x, t) p(t) dt$ and $f(s) = \int_X f(s | x) q(x) dx$.¹⁵ Messages then correspond to lotteries over posterior reputations,

⁹ We may think of the evaluator as being rewarded for predicting as accurately as possible the ability t of the expert based on all the information (m, x) available. It is implicitly assumed that the evaluator is sequentially rational.

¹⁰ In Brandenburger and Polak's (1996) model, the stock market performs a similar role of evaluation but assesses the investment's profitability rather than the quality of the expert's information. In addition, in their model the market observes the message sent by the expert but not the state realization.

¹¹ Zitzewitz (2001) proposes an alternative model in which the market evaluates the quality of the information contained in the forecast with a simple econometric technique, rather than via Bayesian updating.

¹² The payoff to the sender depends entirely on the receiver's belief and may be intangible. The payoff is tangible if it derives from the value of the services provided in a future second and last period by the expert, as in Holmström (1999). Truthful revelation is an equilibrium in this second period.

¹³ This is a psychological game in which the sender's payoff depends on the belief of the receiver. In line with Geanakoplos, Pearce, and Stacchetti (1989), we assume that payoffs have an expected utility formulation. A similar approach has been followed by Bernheim (1994) to measure people's preference for esteem.

¹⁴ As first noticed by Seidmann (1990) in cheap-talk games with intertype agreement, information can nevertheless be transmitted in equilibrium provided that the receiver's decision is based on some additional information (here, the realized state).

¹⁵ This linearity property of the reputational payoff is satisfied in the first of a two-period model, provided that information about ability has no value in the second period. We refer to Li (2004) for an analysis of a reputational cheap-talk game in which this linearity property is violated.

with payoff corresponding to state realizations. Depending on the evaluator’s rule for calculating the posterior reputation, different messages may induce lotteries that are differently appealing to experts with different signals.

An expert strategy φ constitutes a perfect Bayesian Nash equilibrium of the game if for almost all $s \in S$, m solves $\max_{m' \in M} V(m' | s)$ for $\varphi(\cdot | s)$ —almost all $m \in M$, where the value function is computed as just described from the correct conjecture $\hat{\varphi} = \varphi$. Our equilibrium analysis focuses on the sender’s optimal choice among messages sent on the equilibrium path.¹⁶

Throughout the article we will illustrate our results with examples belonging to either of the following two classes of information structures: (i) linear experiment and (ii) location experiment.

Example 1: linear experiment. In this class of examples, the generalized probability density function of the signal conditional on state x and ability t is linear in $t \in T \subseteq [0, 1]$,

$$f(s | x, t) = tg(s | x) + (1 - t)h(s). \tag{3}$$

A signal generated from this experiment can be interpreted as a mixture between an informative and an uninformative experiment.¹⁷ Better experts are more likely to receive a signal drawn from the informative $g(s | x)$ rather than the uninformative $h(s)$. In fact, a more talented expert receives better information in the sense of Blackwell. To see this, consider the garbling of s into \tilde{s} whereby $\tilde{s} = s$ with probability $\tau < 1$, and otherwise \tilde{s} is independently redrawn from $h(s)$. Then $f(\tilde{s} | x, t) = \tau f(s | x, t) + (1 - \tau)h(\tilde{s}) = f(\tilde{s} | x, \tau t)$, so that the garbled signal to an expert of ability $t > 0$ is distributed as the signal to an expert of ability $\tau t < t$.

In the article we make repeated use of three special cases of the linear experiment: (1a) binary experiment, (1b) dichotomous experiment, and (1c) multiplicative linear experiment.

Example 2: location experiment. We consider continuous symmetric location experiments with conditional p.d.f. $f(s | x, t) = g(|s - x| | t)$, where $S = X$ is a normed space, $T = \mathbb{R}$, and g satisfies the following two assumptions: (i) unimodality: $g(\varepsilon | t)$ is a decreasing function of ε , and (ii) monotone likelihood ratio property (MLRP): $g(\varepsilon | t)/g(\varepsilon | t')$ is strictly decreasing in ε when $t' < t$. The latter condition means that a signal realization s closer to x is better news for t . It is simple to see that f has these symmetry and unimodality properties (around x): $f(s | x, t)$ depends on s and x only through $|s - x|$, and it is a decreasing function of $|s - x|$. Clearly, $f(s | x) = g(|s - x|) = \int_T g(|s - x| | t)p(t) dt$ inherits these same properties for any prior $p(t)$. In the interesting special case where $g(\varepsilon | t) = \tilde{g}(\varepsilon/t)/t$ for some fixed p.d.f. \tilde{g} defined on the positive real axis, Lehmann (1988) showed (without assuming symmetry) that the informativeness of the experiment is increasing in the precision t , with respect to monotone decision problems.

Below we will focus on three special cases: (2a) uniform location experiment, (2b) location experiment with prior information, and (2c) normal experiment.

3. Generic impossibility of truthtelling

■ In order to place our contribution, we begin by revisiting known results for the simplest example of reputational cheap talk with two signals.

Example 1a: binary experiment. This very simple linear experiment features two signals $S = \{-1, 1\}$, two states $X = \{-1, 1\}$, two ability types $T = \{t_b, t_g\} \subseteq [0, 1]$ with $t_b < t_g$, and distributions $h(s) = 1/2$ and $g(s | x) = (1 + sx)/2$, so that $f(1 | 1, t) = \Pr(1 | 1, t) = (1 + t)/2 = \Pr(-1 | -1, t) = f(-1 | -1, t)$. Denote the expected ability $Et = p(t_g)t_g + p(t_b)t_b$. This example allows us to gain some initial intuition and introduce the general themes presented in the article.¹⁸

¹⁶ It is assumed that any off-path message is interpreted by the receiver as equivalent to some particular on-path message. Since the number of equilibria is typically small, we do not consider belief refinements.

¹⁷ Note the similarity with Green and Stokey’s (1980) success-enhancing model. In the success-enhancing model the experiment fails with positive probability, in which case the signal is uninformative about the state. In our linear model, the experimenter instead only knows the probability that the experiment is contaminated.

¹⁸ This information structure has been used by Scharfstein and Stein (1990) in the first formulation of reputational cheap talk and further analyzed by Ottaviani and Sørensen (2001).

In a pooling (also known as babbling or uninformative) equilibrium the sender of both types $s = 1$ and $s' = -1$ adopt the same strategy, $\varphi(m | 1) = \varphi(m | -1)$ for all messages m . Given the receiver's correct conjecture of this strategy, $\hat{f}(m | x, t) = \hat{\varphi}(m | s) = \hat{\varphi}(m | s')$ is independent of (x, t) . This implies that the receiver cannot draw any meaningful inference, so that $p(t_g | m, x) = p(t_g)$ for any (m, x) . But then the sender is indifferent among all messages. Note that such a pooling equilibrium always exists.

In a separating (also known as truthtelling, fully revealing, or informative) equilibrium the expert sends a different message depending on the signal received, so that $m(s) = s$ once the message is identified with its meaning. The posterior on ability then satisfies $p(t_g | m, x = m) > p(t_g) > p(t_g | m, x \neq m)$, so that the reputation is updated favorably when the expert is right, and unfavorably when wrong. Since $v(t)$ is increasing and there are two ability types, the sender's objective is to maximize the probability of being perceived to be of type t_g . For truthtelling to be an equilibrium, it must be that the expert with signal s prefers to send $m = s$ rather than $m' = s' \neq s$:

$$q(1 | s)p(t_g | s, 1) + q(-1 | s)p(t_g | s, -1) \geq q(1 | s)p(t_g | s', 1) + q(-1 | s)p(t_g | s', -1) \quad (4)$$

for $s, s' \in \{-1, 1\}$. By Bayes' rule, these two truthtelling constraints are equivalent to $(1 - Et)/2 \leq q(1) \leq (1 + Et)/2$. The truthtelling constraint for the low signal $s = -1$ is violated if the prior belief is biased enough in favor of the high state, $q(1) > (1 + Et)/2$. This is because in this case the expert maximizes the chance of being right by pretending to possess a high signal. More generally, there is an incentive to deviate from a putative separating equilibrium when the prior on the state is biased enough in either direction.

This simple binary example shows that the expert's reputational concern can be either fully compatible or completely incompatible with truthful information revelation. If the prior on the state is intermediate, truthtelling can be sustained in equilibrium. If instead the prior on the state is extreme, no information can be credibly transmitted.

As shown below, full revelation is an artifact of the discrete signal structure, and holds because in a discrete setting the incentive-compatibility constraints required for pure-strategy equilibrium hold with some slack. Turning to a more natural continuous setting, we show that under reasonable conditions on the information structure there are incentives to deviate at the margin from truthful reporting.

We begin by checking whether it is possible to sustain truthtelling in equilibrium. In this cheap-talk environment, the sender's incentives are driven by the receiver's understanding of the meaning of the messages sent, regardless of the language used. For the purpose of checking the possibility of full revelation, it is therefore without loss of generality to restrict the sender to send truthful signals.

By definition, truthful information transmission occurs when $M = S$ and the message sent equals the signal received, so that $\varphi(s | s) = 1$ or, equivalently, $m(s) = s$. Assume for the moment that the receiver naively believes that the sender is applying this truthful strategy, so that $\hat{f}(m | x, t) = f(m | x, t)$. Is truthtelling then the optimal strategy for the sender? If so, truthtelling is an equilibrium. We now show that equilibrium truthtelling is impossible under natural informational assumptions with generic choices of the prior belief and utility function, as also independently observed by Campbell (1998) in a more special case. By ruling out the existence of a truthtelling equilibrium, we conclude that our cheap-talk game cannot have any fully revealing equilibrium.

We say that local truthtelling in the open ball $B \subset S$ is possible, if when the receiver anticipates local truthtelling ($\hat{f}(m | x, t) = f(m | x, t)$ for all $m \in B$) then $V(s | s) = \max_{m \in B} V(m | s)$ for all $s \in B$. In words, there is a whole ball where truthtelling by the sender is a best response to the receiver's anticipation of this. Local truthtelling immediately implies the first-order condition

$$V_m(s | s) = 0 \quad (5)$$

for all $s \in B$.

The experiment is defined to be locally uninformative about ability in the open ball $B \subset S$ if there exist functions $K(t)$ and $\kappa(s | x)$ such that $f(s | x, t) = K(t)\kappa(s | x)$ for almost all $s \in B$, almost all $x \in X$, and almost all $t \in T$. This states that the conditional p.d.f. is separable in the observable outcome (s, x) and the unobservable ability t about which inference is made. The condition implies that the evaluator cannot use the pair (s, x) to make any discriminatory inference on t . Namely, for any two pairs (s, x) and (s', x') with $s, s' \in B$, we have $p(t | s, x) = p(t | s', x')$.

Since the ability type t parameterizes the relation between s and x , local uninformative is a degenerate property that violates the spirit of our model. Local uninformative immediately permits local truthtelling, since the posterior reputation is entirely independent of the message sent. The following converse result is of far greater interest.

Proposition 1 (no truthtelling). Assume that S is a closed, convex subset of \mathbb{R}^J , and that X is a closed subset of \mathbb{R}^L . Assume that $f(s | x, t)$ is bounded and twice continuously differentiable in s , with f and f_s jointly continuous in (x, t) and that the signal structure is not locally uninformative in the open ball $B \subset S$. Local truthtelling in B is then impossible for an open and dense set of prior beliefs $q(x)$ and utility functions $v(t)$.

To understand this result, note that if the signal is not locally uninformative, different message and state pairs (m, x) imply different posterior reputations $p(t | m, x)$. Perturbing the utility function $v(t)$ if necessary, we guarantee that different posterior reputations yield different reputational payoffs $W(m | x)$. The sender is uncertain about the location of x , but perturbations of the prior $q(x)$ finally guarantee that $V_m(s | s)$ varies with s .

Intuitively, the signal s contains information about t , and the expert has an incentive to misrepresent the true s in order to give the most favorable impression. Proposition 1 shows that this incentive is generically embedded in the model. The dependence of the signal on ability generates an incentive to manipulate the report, but this is partly mitigated by the fact that the evaluator cross-checks the report with the realized state.

As shown above, the expert cannot completely reveal a signal s that carries information about ability t . But a signal may in some dimensions carry information solely about the state x , as when the multidimensional signal is $s = (s_1, s_2) \in \mathbb{R}^{J_1} \times \mathbb{R}^{J_2}$ and the signal p.d.f. can be written as $f(s | x, t) = f_1(s_1 | x)f_2(s_2 | x, t)$. In this case, it is easy to see that there is an equilibrium in which s_1 is reported truthfully but no information about s_2 is transmitted, so that the reputation is never updated and the sender is indifferent over all the equilibrium messages. However, since s_2 contains information about (x, t) , there might also be an equilibrium where messages about s_2 are (imperfectly) informative. But since the report of s_1 results in a change in the prior belief on s_2 and thus on the information about s_2 transmitted, truthful reporting of s_1 disappears. The joint formulation $f(s | x, t)$ is therefore most suitable for a complete analysis of the model.

The truthtelling condition (5) also suggests how to use explicit monetary incentives to reinstate truthtelling. If the message sent were verifiable and explicit incentives were allowed, truthtelling could be obtained by offering the reward schedule $R(m) = \int_{-\infty}^m V_m(\tilde{m} | \tilde{m}) d\tilde{m}$ to the expert. Correspondingly, the *ex ante* cost of implementing truthtelling would be $\int_S R(s)f(s) ds$. Notice that the cost could be lower if the reward were allowed to depend also on the realization of the state. For the rest of this article we exclude the possibility of monetary incentives.

Truthtelling is an equilibrium in two very special situations, when the private signal is infinitely more informative about the location of the state (dichotomous experiment) or when there is no public information (uniform location experiment).

Example 1b: dichotomous experiment. In this specification of the linear experiment, the prior on the state $q(x)$ is atomless and the signal is drawn from $f(s | x, t) = t\delta_x(s) + (1 - t)h(s)$, where $\delta_x(s)$ is the Dirac delta function and $h(s)$ is atomless, $T = [0, 1]$ and $X = S \subseteq \mathbb{R}^J$. The expert receives perfect information ($s = x$) with probability t , and otherwise receives an uninformative draw from an atomless distribution.

With this information structure, conditional on signal s , the posterior belief on the state has an atom at $x = s$ and a continuous density over all other states. Moreover, the evaluator that

receives $m = x$ concludes that the signal was derived from the perfectly informative distribution rather than the uninformative one, and that this is good news about the type. Conversely, $m \neq x$ is bad news. Thus, truthful reporting of the signal $m(s) = s$ constitutes an equilibrium, since any other signal has probability zero of turning out to be correct.

Proposition 2 (truthtelling in dichotomous experiment). Truthtelling is an equilibrium in the dichotomous experiment.

Clearly, truthtelling is an equilibrium in the degenerate case with a perfectly informative signal. More generally, truthtelling results when a signal indicates that a state is infinitely more likely than all the other ones.

Example 2a: uniform location experiment. In this specification of the location experiment, there is no prior information on the state. In order to have a proper uniform prior in this location model, the space X should be compact, and complete symmetry is then most easily defined on a symmetric set. We assume that the spaces X and S are both the unit circle, corresponding to the circumference of the unit ball in \mathbb{R}^2 . A real number z indicates a point on the circle in the usual way, giving the anticlockwise distance along the circumference from $(1, 0)$, the circle's origin in the plane.¹⁹

Due to the absence of any prior information about the state, there is no incentive to deviate from truthtelling.

Proposition 3 (truthtelling in uniform location experiment). Truthtelling is an equilibrium in the uniform location experiment.

This result crucially depends on the uniform prior on the state. Truthfully reporting $m = s$ is then equivalent to reporting the mode of the symmetric posterior distribution $q(x | s)$. Since a signal s closer to the state x indicates a higher ability t by the MLRP and the state is concentrated around s , it is advantageous for the sender to send $m = s$ when the receiver interprets m as s . Truthtelling would instead be incompatible with equilibrium for any location experiment with a (generic) proper prior belief on the state.

4. Optimal deviation from truthtelling

■ As we have seen in the previous section, the incentives to deviate destroy any fully revealing cheap-talk equilibrium. In our reputational cheap talk we have derived endogenously the incentive to deviate from truthtelling, depending on the exogenously specified signal structure and value function. In Crawford and Sobel's (1982) model of partisan cheap talk, the incentive to deviate from truthtelling was instead assumed directly as a discrepancy between the sender's and receiver's optimal actions.²⁰

In our reputational model, there is no set direction in which the expert wishes to deviate. In this section we characterize how this direction depends on the primitives of the model. This characterization is important for at least three reasons. First, understanding the pressure to deviate from honesty provides intuition for the impossibility of truthtelling and sheds light on out-of-equilibrium forces. Second, these deviation incentives can explain real-world outcomes of communication, which result when the evaluator is not fully rational.²¹ Third, these incentives persist in the signalling equilibrium that results when the decision-making power is at least partially allocated to the expert.²²

¹⁹ For instance, the numbers $-2\pi, 0, 2\pi$ all indicate $(1, 0)$, while $\pi/2$ indicates $(0, 1)$ in the plane.

²⁰ For instance, in their Lemma 1 they assume that the ideal actions of the sender and the receiver are different for any given state.

²¹ Indeed, existing experimental evidence on simpler cheap-talk games of partisan advice confirms that the deviation incentives influence the outcome even when the players have substantial experience with the game (see Cai and Wang, forthcoming).

²² This is the case in Prendergast and Stole's (1996) reputational signalling model, in which the expert's payoff depends *directly* on the message sent (because of delegation), as well as indirectly through the effect on the updated reputation. This difference explains why their equilibrium is fully revealing, while ours is coarse.

Expanding the first-order condition (5), we now characterize in more detail the forces that drive the sender to bias the report, when honesty is expected.²³

Proposition 4 (deviation incentives). Assume that S is a convex subset of R^J , and that $f(s | t, x)$ is differentiable. If the receiver conjectures truthtelling $m(s) = s$, the marginal incentive to deviate is

$$V_{m_j}(s | s) = \text{cov} \left[v(t), \frac{p_{m_j}(t | s, x)}{p(t | s, x)} \mid s \right] = \text{cov} \left[v(t), \frac{f_{m_j}(s | t, x)}{f(s | t, x)} - \frac{f_{m_j}(s | x)}{f(s | x)} \mid s \right]. \quad (6)$$

Notice that the expectation is taken over the pair (x, t) . The expert is tempted to increase m_j at the margin if this systematically affects the receiver’s posterior likelihood ratios in the direction of greater weight on higher ability. The higher this covariance, the stronger the incentive to deviate. It is naturally in the sender’s interest to convey the impression that the signal s has a larger ratio $f(s | t, x)/f(s | x)$ for higher values of t , because this indicates to the receiver that the sender is likely to be of high ability. Precisely, a local deviation is desirable when it marginally increases the covariance of $d \log(f(s | t, x)/f(s | x))/dm_j$ with the increasing function $v(t)$.²⁴

In the following two examples the direction of the optimal deviation from truthtelling can be easily derived.

Example 1c: multiplicative linear experiment. This experiment generalizes the binary experiment to a continuous setting, with $S = X = [-1, 1]$, $T = [0, 1]$, $h(s) = 1/2$, and $g(s | x) = (1 + sx)/2$ so that $f(s | x, t) = (1 + sxt)/2$.

Plugging the expression of the signal density into equation (6) we obtain

$$V_m(s | s) = [\text{cov}[v(t), t]/2f(s)] \int_{-1}^1 xq(x)/(1 + sxEt) dx.$$

Ignoring the positive leading constant $\text{cov}[v(t), t]/2f(s)$, the optimal deviation direction at s is seen to be determined by the expectation of $x/(1 + sxEt)$. This expectation is decreasing in s , so if it equals zero at some signal, for all other signals there is an incentive to deviate in the direction of that signal. In the special case with prior q symmetric around state $x = 0$, we have $V_m(0 | 0) = 0$, so that the bias is toward the *ex ante* expected state. When instead the prior belief is strongly biased in either direction, Ottaviani and Sørensen (2006) have shown that there is an incentive to deviate in that same direction for all signals, again reflecting a bias toward the expected.

Example 2b: location experiment with prior information. In this example, the signal is drawn from a location model where $S = X = \mathbb{R}$, and *a priori* x has a strictly quasi-concave distribution, symmetric around the prior location μ . Thus $q(x) = h(|x - \mu|)$ where h is a strictly decreasing function.

The expert always exhibits a bias toward the *ex ante* expected state μ .²⁵

Proposition 5 (conservative bias in location experiment). In the location experiment with prior information, there is a conservative deviation incentive: $V_m(s | s) > 0$ when $s < \mu$, and $V_m(s | s) < 0$ when $s > \mu$.

²³ This covariance characterization is similar to the one obtained by Dewatripont, Jewitt, and Tirole’s (1999) for a general version of Holmström’s (1999) career concern problem with hidden action. In our hidden information context without commitment instead, this characterization relates to deviation rather than equilibrium incentives.

²⁴ The value function $v(t)$ is a strictly increasing transformation of t , so a simple change of variables in the statistical model (from t to v) here allows us to assume without loss of generality that $v(t) = t$.

²⁵ Ottaviani and Sørensen (forthcoming) use a special case of this model where the prior $q(x)$ and the sender’s posterior $q(x | s)$ are both normal. They establish that it is optimal to deviate from truthful reporting of s to pretending that the signal is $\bar{s} = E[x | s]$.

5. Equilibrium predictions

- In the following examples, equilibria have a partition structure.

Example 1a (continued). When $q(1) \in [(1 - Et)/2, (1 + Et)/2]$, the two truth-telling constraints (4) for $s, s' \in \{-1, 1\}$ are satisfied, so that there is a separating equilibrium. Under this condition, one can also construct a hybrid mixed-strategy equilibrium, featuring elements of both separating and pooling. For $1/2 \leq q \leq (1 + Et)/2$, the expert with signal $s = 1$ sends message $m = 1$, while the expert with signal $s = -1$ sends the low message with probability $\varphi(m = -1 | s = -1) = (2q - 1)/(q - (1 - Et)/2) \in [0, 1]$ and the high message $m = 1$ with complementary probability. The message thus garbles the signal. By symmetry, for $(1 - Et)/2 \leq q \leq 1/2$, there is a hybrid equilibrium with a similar structure. In conclusion, for $q(1) \in [(1 - Et)/2, (1 + Et)/2]$, there are three equilibria: separating, pooling, and hybrid; for $q \notin [(1 - Et)/2, (1 + Et)/2]$, the only equilibrium outcome is pooling. Nevertheless, in this example the most informative equilibrium is unique.

Example 1c (continued). Consider a special version of the multiplicative linear experiment with two states $X = \{-1, 1\}$ showcased by Ottaviani and Sørensen (2006). When the prior $q(1)$ is close enough to $1/2$, there are two informative binary equilibria, and the posterior belief of the indifferent type is necessarily more extreme than the prior belief. This means that there is a region of types of sender with a posterior belief between $1/2$ and q who nevertheless report an inclination toward the *ex ante* unlikely state. In this sense, there is contrarian behavior in this model. In addition, notice that even in this simple example there is no unique most-informative equilibrium. If the prior $q(1)$ is far enough from $1/2$, the expert always wishes to bias her report in the same direction and is then unable to communicate any information in equilibrium.

Example 2b (continued). In the symmetric location experiment with normal prior $x \sim N(\mu, 1/\nu)$, Ottaviani and Sørensen (forthcoming) have verified that there always exists a two-message equilibrium of the following kind. The sender reports whether the signal is above or below the prior mean μ . When observing a signal $s > \mu$, the sender expects high values of the state to be realized that are more likely to result in favorable reputational updating when the sender is believed to have received signal above rather than below the prior mean. This is then an equilibrium, since by symmetry the sender with $s = \mu$ is indifferent among these two messages.

For the application of this model to the predictions of professional experts, we need to discuss how information is communicated. In equilibrium, the receiver understands that signals are garbled into the sender's message m . The advice given by the expert is typically used by a decision maker, whose beliefs $f(x | m)$ are unambiguously determined in equilibrium. The model predicts that in equilibrium the resulting belief $f(x | m)$ is unbiased, being derived from Bayesian updating, but is less informative than the expert's private belief $f(x | s)$ in the sense of Blackwell.

In the rest of this section, we discuss three difficulties encountered when attempting to test directly this equilibrium prediction. The first difficulty is that equilibrium beliefs are often unobservable. Nevertheless, in a number of settings, such as forecasting, the sender's reports are observable. It is then natural to use these reports for the purpose of testing the theory.

When using the sender's reports, the empirical comparison of the experts' literal statements with the outcome of the predicted variable x is impossible. This is because the language used to send equilibrium messages is indeterminate, as in any cheap-talk game. To overcome this second difficulty, it is therefore necessary to make an assumption about the language used in equilibrium. The natural language in this setting dictates that the expert communicates the equilibrium posterior belief on the state (or its mean $E[x | m]$) or recommends the corresponding optimal action conditional on that belief. If so, the message m is translated into the best predictor on the state incorporating all its informational content. Statements in this language can be easily compared with

the realized state. The empirically observed report would be more inaccurate than the conditional expectation obtained from the expert's signal.²⁶

A direct test of our theory would then be based on the regression

$$x = \alpha_0 + \alpha_1 m + \alpha_2 y + \alpha_3 s + \varepsilon, \quad (7)$$

where x is the realized state, m the report, y any publicly known variable, s the private information of the expert, and ε the error term. Unbiasedness requires that when y and s are excluded, the remaining coefficients are restricted to $\alpha_0 = 0$ and $\alpha_1 = 1$. Efficiency requires that all information available to the expert has no additional predictive power in the regression, i.e., $\alpha_2 = \alpha_3 = 0$. Identifying m with the prediction on the state $E[x \mid m, y]$, our reputational cheap-talk model predicts unbiasedness and efficiency only with respect to public information $\alpha_2 = 0$. According to our coarseness result, the message sent is not a sufficient statistic for the expert's private information. Furthermore, it is easy to show that the MLRP of s, x implies the MLRP of s, x conditional on any realization m , when s is a Blackwell-sufficient experiment for m (see Ottaviani and Prat, 2001). Thus our model predicts that $\alpha_3 > 0$. Direct test of this prediction would require access to the expert's private information, but this is rarely available. This is the third difficulty encountered when using this direct-testing strategy.

This difficulty has led empiricists to test for the orthogonality property of reports, i.e., the fact that the report is uncorrelated with its error. In the case of forecasting, orthogonality is tested by regressing the realized forecast error on the forecasts. For further discussion of the predictions of the theory when applied to strategic forecasting we refer to Ottaviani and Sørensen (forthcoming), where we compare the predictions of the reputational cheap-talk theory to those of the forecasting-contest theory and extend both theories in a number of directions relevant for empirical testing.

The equilibrium loss of information typically results in a welfare loss for the current decision maker. In addition, future employers of the expert will have less information about the expert's true ability. Future decision makers would prefer the signal not to be garbled if they value information about the expert's ability.²⁷ If the utility function $v(t)$ is linear, the sender's *ex ante* expected reputational value of sending any message profile is equal to its prior value. Therefore, the expert is indifferent in *ex ante* terms between the different equilibria. In expectation, no one benefits from the fact that information transmitted in equilibrium is less precise than the information possessed by the expert.

6. Known ability

■ Up to now we have assumed that the expert has no superior information about ability compared to the evaluator. In some situations experts have acquired such information from previous experience. We now investigate the robustness of our findings to cases in which the expert observes signal and ability (s, t) before reporting to the evaluator, as first considered by Trueman (1994).²⁸

□ **Model extension.** An expert of ability t who receives signal s has posterior on the state

$$q(x \mid s, t) = \frac{f(s \mid x, t)}{f(s \mid t)} q(x). \quad (8)$$

²⁶ In case many identical experts are polled simultaneously, their reports should be concentrated on a finite set of positions. See also Section 5.

²⁷ Information about the quality of information received can enable future employers to make better decisions. This information is also valuable in some dynamic settings with competition. For example, Prat (2005) develops two simple two-period settings in which future employers necessarily benefit from information about ability.

²⁸ See also Avery and Chevalier (1999) and Levy (2004).

The evaluator's conjecture of the sender's strategy is now denoted by $\hat{\varphi}(m | s, t)$. The posterior reputation is $p(t | m, x) = \hat{f}(m | x, t)p(t) / \hat{f}(m | x)$, where the conditional chance of message m computed by the evaluator is $\hat{f}(m | x, t) = \int_S \hat{\varphi}(m | s, t)f(s | x, t)ds$. The expected reputational payoff of message m for a sender with signal s and ability t is assumed to be

$$V(m | s, t) \equiv \int_X \int_T v(t')p(t' | m, x)dt'q(x | s, t)dx, \quad (9)$$

for a strictly increasing utility function $v(t)$.²⁹

We begin by revisiting the binary model with known ability.

Example 1a (continued). Suppose that the receiver continues to conjecture that the expert truthfully reveals the signal. The resulting reputation is calculated as before, and therefore still has $p(t | -1, -1) = p(t | 1, 1) > p(t | -1, 1) = p(t | 1, -1)$. When $q(1) > (1 + t)/2$, the expert believes state $x = 1$ to be the more likely outcome and so wishes to deviate from honest reporting of $s = -1$. The conservative bias is thus preserved from the case of unknown ability, although more able experts have less incentive to deviate.

In the range $q(1) \in [(1 + t_b)/2, (1 + t_g)/2]$, a natural conjecture of the receiver is that the expert reports the most likely state. This means that the high type sends a message $m = s$, while the low type always plays $m = 1$. But then message $m = -1$ reveals that the expert is of the high type, and it therefore dominates message $m = 1$. Observe that this deviation incentive is now opposite to the one resulting with unknown ability.

Ottaviani and Sørensen (2001) establish in Lemma 4 that the most informative equilibrium involves two messages with the following properties. For $q(1) \in [1/2, (1 + t_b)/2]$, there is pooling on the ability dimension but separation on the signal dimension, while for $q(1) \in [(1 + t_b)/2, 1]$ both high- and low-ability experts with $s = 1$ send the positive message $m = 1$, the high-ability sender with $s = -1$ always sends the negative message $m = -1$, but the low-ability sender with $s = -1$ strictly randomizes between the positive and negative message. This equilibrium strikes a balance between the two natural conjectures that we considered above. Notice that when $q(1) > (1 + t_g)/2$, any expert who reports $m = -1$ is not conservative, but is in fact adopting a position that is biased away from the most likely state. On the other hand, the low-ability expert with $s = -1$ is driven by the conservatism incentive to sometimes report $m = 1$.

□ **Impossibility of truthtelling.** It is immediate to see that there cannot be a fully revealing equilibrium whereby both the signal realization s and the ability type t are communicated truthfully, $m(s, t) = (s, t)$. Otherwise, each expert would want to claim to have the highest possible ability type. Therefore, there must be some form of pooling in the two-dimensional type space $S \times T$. Clearly, truthful reporting of ability and pooling on the signal, $m(s, t) = t$, is incompatible with equilibrium.

Consider next the strategy of truthfully reporting the signal $m(s, t) = s$ and pooling on the type. By extending Proposition 1, it can be shown that if the adviser knows her own ability, there is no equilibrium with truthful revelation of the signal for an open and dense set of priors $q(x)$ and utility functions $v(t)$.

Examples 1b and 2a (continued). In the dichotomous and uniform location experiments with known ability, it can be shown that there is an equilibrium with truthful revelation of the signal under the assumptions of Proposition 2 and 3 respectively.

□ **Deviation from truthful revelation of signal.** The marginal incentive to deviate from truthful revelation of the signal can be expressed in the style of Proposition 4 as

²⁹ This is a strong assumption in this setting, because an expert with private information about her own ability knows better than the market how her reputation will be updated in later periods. In a fully fledged dynamic model, an expert's prospects of future earnings would then depend on ability.

$$\begin{aligned}
 V_m(s \mid s, t) &= E \left[v(t') \frac{p_m(t' \mid m(s, t), x)}{p(t' \mid s, x)} \frac{q(x \mid s, t)}{q(x \mid s)} \mid s \right] \\
 &= \frac{f(s)}{f(s \mid t)} E \left[v(t') \frac{p_m(t' \mid m(s, t), x)}{p(t' \mid s, x)} \frac{f(s \mid x, t)}{f(s \mid x)} \mid s \right]. \tag{10}
 \end{aligned}$$

Again, the expectation is taken over the pair (x, t') , where t' represents the variable over which the receiver’s beliefs are defined. Equation (10) is similar to (6), only complicated by the presence of a third factor that reflects the sender’s desire to bias the reported signal in a direction associated with x values deemed likely given t .

We now verify that the bias toward the expected is robust in the multiplicative linear experiment and the location experiment with a unimodal prior.

Example 1c (continued). Using criterion (10), we find

$$V_m(s \mid s, t) = [\text{cov}[v(t), t]/2f(s \mid t)] \int_{-1}^1 xq(x)(1 + stx)/(1 + sxEt)^2 dx.$$

When $t = Et$, this reduces to $V_m(s \mid s)$ from the case of unknown ability, so this average sender has precisely the same incentive to deviate as before. For any t , the sign of $V_m(s \mid s, t)$ is dictated by the integral $\int_{-1}^1 xq(x)(1 + stx)/(1 + sxEt)^2 dx$, which has a derivative in t of the same sign as s . Thus, for any $s > 0$, $V_m(s \mid s, t)$ is increasing in t . In the special case with prior q symmetric around state 0, we have seen in the unknown-ability version of the model that $V_m(0 \mid 0) = 0$ and $V_m(s \mid s)$ is of opposite sign to s . We then conclude that in this case the bias toward the expected is weaker for more able experts.

Example 2b (continued). In the location experiment with prior information and known ability, it can be shown (as in Proposition 5) that there is a conservative deviation incentive: $V_m(s \mid s, t) > 0$ when $s < \mu$, and $V_m(s \mid s, t) < 0$ when $s > \mu$.

□ **Deviation from truthful revelation of posterior.** Having shown that truthful reporting of the signal is impossible, we now investigate truthful reporting of the posterior belief $q(x \mid s, t)$ —this belief often pools different combinations of signal s and ability t . Note that higher-ability experts more often have extreme posterior beliefs about the state. Intuitively, there should then be an incentive to send more extreme messages in order to indicate ability. This intuition is verified in Example 1a as well as in the following two examples, where the previous bias toward the expected is replaced by a bias away from the expected.³⁰

Example 1c (continued). In this model, the sender’s posterior $q(x \mid s, t)$ depends on $(s, t) \in [-1, 1] \times [0, 1]$ only through the product st . Obviously, truthful reporting of the posterior is impossible, because deviation to messages $m \in \{-1, 1\}$ gives the highest possible reputation, $t = 1$ for sure. Through $\hat{f}(m \mid x, t) = (1 + mx)/2t$ for $|m| \leq t$ and $p(t \mid m, x) = p(t)/[2t \int_{|m|}^1 p(t')/2t' dt']$ for $t \geq |m|$, we obtain

$$V_m(m(s, t) \mid s, t) = \frac{p(m)}{2m} \frac{\int_{|m|}^1 (v(t') - v(|m|))p(t')/2t' dt'}{(\int_{|m|}^1 p(t')/2t' dt')^2},$$

which has the same sign as m . The deviation incentive is therefore now away from the middle messages.

Example 2c: normal location experiment. The normal location experiment is a special version of

³⁰ This exaggeration incentive persists in equilibrium in Prendergast and Stole’s (1996) model, where the decision is delegated to the sender and the evaluator does not observe the state realization. This incentive also persists in equilibrium in Zitzewitz’s (2001) model, where the evaluation of forecast quality is performed with a realistic econometric technique.

the location experiment with symmetric prior, in which the prior is $x \sim N(\mu, 1/\nu)$ and the signal is $s \mid x, t \sim N(x, 1/t)$.³¹

In this example, the sender's posterior is $x \sim N((ts + \nu\mu)/(t + \nu), 1/(t + \nu))$ and this distribution identifies the pair (s, t) .³² A natural candidate strategy in the normal location experiment is to report the posterior mean, $m(s, t) = (ts + \nu\mu)/(t + \nu)$. We can prove that in this setting there is again an incentive to exaggerate.

Proposition 6 (exaggeration in normal experiment). In the normal experiment, when the receiver conjectures truthful reporting of the posterior mean, there is an incentive to exaggerate as $V_m(m(s, t) \mid s, t) < 0$ when $m < \mu$ and $V_m(m(s, t) \mid s, t) > 0$ when $m > \mu$.

□ **Equilibrium predictions.** We have argued that deviation incentives may depend on whether the ability is known by the sender. But our deviation incentives are confounded in equilibrium, as concluded in Example 1a. We now elaborate on this point.

Example 1c (continued). Consider the multiplicative linear experiment with two states $X = \{-1, 1\}$. We investigate the existence of equilibria in which the threshold of indifference between a message and another are iso-posterior curves. In a binary equilibrium with threshold value k of st , message m is sent for $-1 \leq st \leq k$ and message m' is sent for $k < st \leq 1$. Unlike in the unknown-ability case, in this example (as well as in Example 1a), for any prior belief there exist informative equilibria, where experts report whether $E[x \mid s, t]$ is above or below a certain threshold value.

Proposition 7 (informative equilibrium with known ability). In the multiplicative linear experiment with known ability, a binary equilibrium exists for any prior $q(x)$.

To understand this result, note that a message sent exclusively by the highest possible type gives the highest possible reputation, regardless of the realized state of the world. No matter how extreme is the prior belief on the state, the strongest types always have the self-confidence required to send a message opposite to the prior. If a sufficiently small set of good types are sending the message, they signal that they are good and secure a good minimum reputation, even when the state of the world turns out against them.

Example 2b (continued). In the symmetric location experiment, the binary equilibrium reported in Section 5 also exists when the expert knows her own known ability. If the signal is very far from μ , it is then conservative behavior to send a message that pools with signals predominantly closer to μ . Conversely, when the signal is very close to μ , it is an exaggeration to pool with signals mostly farther from μ . Since more able experts have more variable posterior beliefs on the location, they are more likely to exhibit conservative behavior in this equilibrium.

Although we can say something meaningful about the direction of deviations in these two examples, the only general conclusion about the equilibrium of the model is that biases cannot be unidirectional. Consider any message m sent by the sender whenever $(s, t) \in M$ where M is some subset of $S \times T$, and focus on the implied message of $E[x \mid M]$. This generally differs from the expert's own $E[x \mid s, t]$, since the expert has finer information. Now, the average of the deviations $E[x \mid M] - E[x \mid s, t]$ is necessarily equal to zero, since by the law of iterated expectations we have $E[x \mid M] = E[E[x \mid s, t] \mid M]$. Thus a rational receiver is not fooled in equilibrium. While some experts may be biased in one direction when reporting M , other experts reporting the same message M must necessarily have the opposite bias.

Our examples suggest that the equilibria tend to select more able experts to be more frequently biased toward the expected. This effect arises because they end up pooling with lower-ability experts, whose posterior beliefs are less variable.

³¹ Notice that this normal experiment is different from that used in Ottaviani and Sørensen (forthcoming), in which $q(x \mid s)$ rather than $q(x \mid s, t)$ is normally distributed.

³² More general location experiments have often only trivial iso-posterior sets corresponding to single ability levels, which clearly cannot be truthfully reported.

7. Multiple experts and relative reputational concerns

[A] decision was made at the outset of this study not to disclose the sources of the forecasts evaluated . . . forecasters are rivals or competitors . . . Any statement bearing on the relative quality of a forecaster's product could be used in this competition. (Zarnowitz, 1967, p. 1)

■ Can competition between experts affect the amount of information credibly communicated? For instance, full information revelation results in equilibrium when consulting simultaneously multiple perfectly informed experts in the partisan cheap-talk model of Crawford and Sobel (1982), as shown by Krishna and Morgan (2000) and Battaglini (2002). Consider instead multiple professional experts with conditionally independent signals. If they simultaneously report their messages and care only about their own *absolute* reputation, the equilibrium is the same as in the single-expert model.

If instead the market rewards those with better reputation more if they are scarcer, experts should care about their *relative* reputation. One could expect that more differentiation and perhaps more information revelation would result in the presence of relative reputational concerns. We now show that this is not the case when reputational preferences have a von Neumann-Morgenstern representation and experts have conditionally independent signals.

We consider simultaneous reporting by experts $i = 1, \dots, N$. Expert i 's von Neumann-Morgenstern payoff $u^i(t^i, t^{-i})$ is assumed to depend also on the ability of all other experts $t^{-i} \equiv (t^1, \dots, t^{i-1}, t^{i+1}, \dots, t^N)$, with u^i increasing in t^i but possibly decreasing in t^j for $j \neq i$. We also make the natural assumptions that the state of the world and the ability types of the experts are independently distributed and that the experts' signals are independent conditional on state and ability.

The independence assumption implies stochastic independence of posterior reputations of different experts updated after the reports and observation of the state of the world. Moreover, only an expert's own message (and the state of the world) influences the updating of the reputation of that expert. According to the Martingale property of updated Bayesian beliefs, the expected posterior reputations of other experts equal the prior reputations. Finally, the von Neumann-Morgenstern payoff is linear in those beliefs. Thus we have the next general result.

Proposition 8 (irrelevance of relative reputation). Assume that the experts have von Neumann-Morgenstern payoffs, and that their signals are independent conditional on state and ability. In equilibrium of the relative reputation model with unknown and known own ability and any experiment $f(s | x, t)$, expert i behaves as in the absolute reputation model with increasing utility function $v^i(t^i) = E_{t^{-i}}[u^i(t^i, t^{-i})]$.

To generate new and interesting results, a relative-reputations model must then either assume that there is correlation of experts' signals conditional on the state and ability draw or give up the von Neumann-Morgenstern formulation. For an investigation of relative reputational concerns in a binary model with conditionally correlated signals, see Effinger and Polborn (2001).

It is worth remarking that our von Neumann-Morgenstern formulation is rather restrictive in this setting with multiple experts and does not allow for the market to reward experts on the basis of a comparison of some summary statistics of their updated reputation. For instance, the case in which the expert with highest expected ability receives all the rewards cannot be modelled with von Neumann-Morgenstern payoffs.

8. Conclusion

■ We conclude by briefly summarizing the answers to the five questions addressed in this article and by suggesting some avenues for future research.

First, when is the expert's concern for accuracy compatible with truthful information transmission? The main result of this article is that truthtelling is generally not an equilibrium when the signal and the state can be cross-checked to update beliefs about the expert's ability. In a putative fully revealing equilibrium, the signal and the realized state are informative about

ability, giving an incentive to the expert to manipulate the report of the signal in order to generate a better reputation. In contrast with the canonical model of partisan advice, truthtelling in a professional setting is possible, but only under nongeneric conditions. If there is some interval of signals that is truthfully reported in equilibrium for an open and dense set of priors over the state and utility functions, then the signal structure must satisfy the very special property of local unformativeness.

Second, in which direction does the expert wish to bias the report, when believed to be truthful? The expert gains from pretending to have a higher signal than the one actually observed if this affects the signal likelihood ratio in the direction of higher ability. In the examples we have seen that this often implies a bias toward the expected.

Third, what reporting strategy should the investor rationally expect the expert to use in equilibrium? While there is always a pooling equilibrium in which the expert communicates no information, there are often partially separating equilibria in which some information is conveyed. The report garbles the information about the state of the world as well as the expert's ability. In cheap-talk equilibrium, the rational receiver unravels any attempt by the sender to bias the report uniformly in one direction.

Fourth, what is the effect of the expert's knowledge about own ability? This introduces an incentive to send more extreme messages in order to signal ability, but in equilibrium the more informed experts are forced to be more often biased toward the expected.

Fifth, what is the effect of competition among experts? We have shown that relative reputational concerns make no difference to the model's predictions if the payoff has a von Neumann-Morgenstern specification and the signals are conditionally independent.

Finally, the stylized model studied in this article leaves open a number of questions for further research. While the implicit incentives provided by the market discipline the expert's behavior, they also increase the scope for strategic manipulation in the revelation of a given level of information. It is natural to ask how these problems can be overcome with optimally designed explicit incentives. A starting point for investigating the interaction between explicit and implicit incentives is provided by Holmström and Ricart i Costa (1986) for the case in which ability adds instead to the value of the output produced. It would also be interesting to extend our model to allow the expert to become more informed by acquiring costly signals. The point of departure would be Osband's (1989) study of explicit incentives for truthtelling and information acquisition by forecasters in the absence of reputational concerns.

Appendix

■ Proofs of Propositions 1–8 follow.

Proof of Proposition 1. We first specify the topology on the set of prior beliefs and utility functions as the product topology of the following two topologies. On the set of probability distributions over X , employ the weak topology (see Parthasarathy, 1967, p. 40), which lets the net q_λ of measures converge to the limit measure q if and only if $\int v dq_\lambda$ converges to $\int v dq$ for all bounded real-valued continuous functions v . On the set of weakly increasing continuous functions on T , employ the topology of uniform convergence.

The set in which truthtelling holds on B is defined by the collection of weak inequalities $V(s | s) \geq V(m | s)$ for all $m, s \in B$. By continuity of the integrals with respect to the prior belief $q(x)$ and utility function $v(t)$, this set is closed, being the intersection of closed sets. Its open complement is the set of q, v for which local truthtelling in B is impossible.

This open set is also dense, since from any pair q, v it is possible to find another pair q', v' arbitrarily close to q, v such that local truthtelling fails. This is shown analytically by expanding (5). We have

$$V_m(m | s) = \int_X W_m(m | x)q(x | s) dx = \int_X \left[\int_T v(t)p_m(t | m, x) dt \right] q(x | s) dx,$$

so that

$$V_m(s | s) = \int_X \int_T v(t)p_m(t | s, x) dt \frac{f(s | x)}{f(s)} q(x) dx.$$

If truthtelling holds for all $s \in B$ under local perturbations in q , then for almost all $x \in X$, $0 = \int_T v(t)p_m(t | s, x) dt f(s | x)$. Lemma A1 (stated and proven below) guarantees that, unless $p_m(t | s, x) = 0$ for all t , there exists a weakly increasing, continuous $d : T \rightarrow \mathbb{R}$ such that $\int_T d(t)p_m(t | s, x) dt \neq 0$. Reversing this logic, if truthtelling is to be robust against the local perturbations $v + \varepsilon d$ for $\varepsilon > 0$ to v , then for almost all $s \in B$, almost all $x \in X$, and all $t \in T$, $0 = p_m(t | s, x)$. Note that

$$\frac{\partial p(t | s, x)}{\partial m_j} = p(t) \frac{(\partial f(s | x, t)/\partial s_j) f(s | x) - (\partial f(s | x)/\partial s_j) f(s | x, t)}{(f(s | x))^2}.$$

So $0 = p_m(t | s, x)$ implies for almost all $s \in B$, almost all $x \in X$, almost all $t \in T$, and all $j = 1, \dots, J$:

$$\frac{\partial f(s | x, t)/\partial s_j}{f(s | x, t)} = \frac{\partial f(s | x)/\partial s_j}{f(s | x)}.$$

This condition states that the ratios $(\partial f(s | x, t)/\partial s_j)/f(s | x, t)$ do not depend on t . The ratio $(\partial f(s | x, t)/\partial s_j)/f(s | x, t)$ is equal to $d \log(f(s | x, t))/ds_j$, so through integration the ratios determine $\log(f(s | x, t))$ up to an additive constant. Thus we can conclude that there exist functions $K(t)$ and $g(s | x)$ such that for almost all $s \in B$, almost all $x \in X$, and almost all $t \in T$ we have $f(s | x, t) = K(t)g(s | x)$. The signal is then locally uninformative, in violation of the assumption. *Q.E.D.*

Lemma A1. Consider any continuous function $g : T \rightarrow \mathbb{R}$. Either $g(t) = 0$ for almost all t , or there exists a weakly increasing, continuous $d : T \rightarrow \mathbb{R}$ such that $\int_T g(t)d(t) dt \neq 0$.

Proof. Unless $g = 0$ almost everywhere, there exists some $\varepsilon > 0$ and $t \in T$ such that $|g(t')| > |g(t)|/2$ for all $t' \in (t - \varepsilon, t + \varepsilon)$. Let two weakly increasing, continuous functions $d_1, d_2 : T \rightarrow \mathbb{R}$ be defined as follows. Let $d_1(t') = 0$ when $t' \leq t - \varepsilon$, $d_1(t') = t' - t + \varepsilon$ when $t' \in (t - \varepsilon, t)$, and $d_1(t') = \varepsilon$ when $t' \geq t$. Likewise, let $d_2(t') = 0$ when $t' \leq t$, $d_2(t') = t' - t$ when $t' \in (t, t + \varepsilon)$, and $d_2(t') = \varepsilon$ when $t' \geq t + \varepsilon$. Also, define $d_3 = d_1 - d_2$. It is simple to verify that d_3 is continuous, zero outside $(t - \varepsilon, t + \varepsilon)$, strictly positive inside $(t - \varepsilon, t + \varepsilon)$, and that $\int_T g(t)d_3(t) dt \neq 0$. Since $\int_T g(t)d_3(t) dt = \int_T g(t)d_1(t) dt - \int_T g(t)d_2(t) dt$, it follows that $\int_T g(t)d_1(t) dt \neq 0$ or $\int_T g(t)d_2(t) dt \neq 0$, as desired. *Q.E.D.*

Proof of Proposition 2. It is seen immediately that $f(s | x) = (Et)\delta_x(s) + (1 - Et)h(s)$ and $f(s) = (Et)q(s) + (1 - Et)h(s)$. Assuming truthtelling, the posterior reputation is

$$p(t | m, x) = \frac{f(m | x, t)}{f(m | x)} p(t) = \begin{cases} tp(t)/Et & \text{if } x = m \\ (1 - t)p(t)/(1 - Et) & \text{if } x \neq m \end{cases}.$$

The reputation obtained after the realization $x = m$ dominates in the first-order stochastic sense the reputation following $x \neq m$. The posterior on x is

$$q(x | s) = \frac{f(s | x)}{f(s)} q(x) = \frac{(Et)q(s)}{(Et)q(s) + (1 - Et)h(s)} \delta_x(s) + \frac{(1 - Et)h(s)}{(Et)q(s) + (1 - Et)h(s)} q(x),$$

an average between an atom at $x = s$ and the continuous prior $q(x)$.

By sending $m = s$, there is a positive probability that $x = m$ and the reputation will be favorably updated. If instead $m \neq s$, there is probability zero that $x = m$, so that the updating which results is necessarily unfavorable. The sender prefers the chance of a favorable updating and so truthfully reports $m = s$. *Q.E.D.*

Proof of Proposition 3. Since the receiver anticipates truthtelling, the posterior reputation is

$$p(t | m, x) = \frac{p(t)f(m | x, t)}{f(m | x)} = \frac{p(t)g(|m - x| | t)}{g(|m - x|)}.$$

We see that $p(t | m, x)$ depends only on $|m - x|$. By the MLRP, the smaller is $|m - x|$, the better news for t , i.e., $p(t | m, x)$ is better in the first-order stochastic dominance sense. For any increasing v we conclude that $W(m | x) \equiv \int_T v(t)p(t | m, x) dt$ depends only on $|m - x|$, and is a decreasing function of $|m - x| \in [0, \pi]$.

Since $q(x)$ is the uniform distribution, the sender's posterior belief on x is described by the p.d.f. $q(x | s) = f(s | x)/f(s)$. By symmetry and unimodality of $f(s | x)$, this distribution of x is symmetric and unimodal around s . Thus, $q(x | s)$ depends only on $|x - s|$ and is decreasing in $|x - s| \in [0, \pi]$.

We now show that these properties of W and q imply that $m = s$ maximizes $V(m | s) = \int_X W(m | x)q(x | s) dx$ over m , so that truthtelling is optimal. By symmetry, it suffices to consider $m \in [s, s + \pi]$ and prove that $V(s | s) \geq V(m | s)$. Note that half of the space X is closer to s than to m , namely the values of x in the interval $[(s + m - 2\pi)/2, (s + m)/2]$. We have

$$\begin{aligned}
 V(s | s) - V(m | s) &= \int_{(s+m-2\pi)/2}^{(s+m)/2} [W(s | x) - W(m | x)]q(x | s) dx + \int_{(s+m)/2}^{(s+m+2\pi)/2} [W(s | x) - W(m | x)]q(x | s) dx \\
 &= \int_{(s+m-2\pi)/2}^{(s+m)/2} [W(s | x) - W(m | x)]q(x | s) dx \\
 &\quad + \int_{(s+m-2\pi)/2}^{(s+m)/2} [W(s | s+m-x) - W(m | s+m-x)]q(s+m-x | s) dx \\
 &= \int_{(s+m-2\pi)/2}^{(s+m)/2} [W(s | x) - W(m | x)][q(x | s) - q(x | m)] dx,
 \end{aligned}$$

where the first equality is by definition, the second uses the change of variable $y = m + s - x$ in the second integral, and the last follows from W and q depending on their arguments only through their distance. Since $[(s + m - 2\pi)/2, (s + m)/2]$ has x values closer to s than m , we have $W(s | x) \geq W(m | x)$ and $q(x | s) \geq q(x | m)$. Then the integrand is always nonnegative and so the integral is nonnegative, proving $V(s | s) - V(m | s) \geq 0$. *Q.E.D.*

Proof of Proposition 4. Expanding the derivative of $V(m | s)$ with respect to coordinate m_j at the truthful $m = s$, we obtain

$$\begin{aligned}
 V_{m_j}(s | s) &= \int_X \int_T v(t)p_{m_j}(t | s, x) dt q(x | s) dx \\
 &= \int_X \int_T v(t) \frac{p_{m_j}(t | s, x)}{p(t | s, x)} \frac{f(s | x, t)p(t)}{f(s | x)} dt \frac{f(s | x)q(x)}{f(s)} dx \\
 &= \int_X \int_T v(t) \frac{p_{m_j}(t | s, x)}{p(t | s, x)} q(x, t | s) dt dx,
 \end{aligned}$$

where $q(x, t | s)$ denotes the posterior belief on (x, t) , i.e., $f(s | x, t)p(t)q(x)/f(s)$ by Bayes' rule. Since $E[p_{m_j}(t | s, x)/p(t | s, x) | s] = 0$ for all s , we have $V_{m_j}(s | s) = \text{cov}(v(t), p_{m_j}(t | s, x)/p(t | s, x) | s)$. Finally, from $p(t | s, x) = f(s | x, t)p(t)/f(s | x)$, we obtain

$$\frac{p_{m_j}(t | s, x)}{p(t | s, x)} = \frac{f_{m_j}(s | t, x)}{f(s | t, x)} - \frac{f_{m_j}(s | x)}{f(s | x)}.$$

Q.E.D.

Proof of Proposition 5. As in the proof of Proposition 3, $W(m | x)$ is a decreasing function of $|m - x| \in \mathbb{R}_+$. The sender's posterior belief is $q(x | s) = f(s | x)q(x)/f(s) = g(|s - x|)h(|x - \mu|)/f(s)$. Now,

$$\begin{aligned}
 f(s)V_m(s | s) &= f(s) \int_{-\infty}^{\infty} W_m(s | x)q(x | s) dx \\
 &= \int_{-\infty}^s W_m(s | x)g(|s - x|)h(|x - \mu|) dx + \int_s^{\infty} W_m(s | x)g(|s - x|)h(|x - \mu|) dx \\
 &= \int_s^{\infty} W_m(s | x)g(|s - x|)[h(|x - \mu|) - h(|2s - x - \mu|)] dx,
 \end{aligned}$$

where the last equality used the change of variable $y = 2s - x$ with the symmetry properties $W_m(s | x) = -W_m(s | 2s - x)$ and $g(|s - x|) = g(|s - (2s - x)|)$. For any $x > s$ we have $W_m(s | x) < 0$, since W is decreasing in the distance of its arguments. Now suppose $s < \mu$, then $2s - \mu < s$, so every $x > s$ satisfies $|x - \mu| < |2s - x - \mu|$. Since h is decreasing, it follows that the integrand is positive for almost all x , and so the integral is positive. Symmetrically, when $s > \mu$ the integral is negative. *Q.E.D.*

Proof of Proposition 6. The message m is a linear transformation of s , so that $m | x, t \sim N(m(x, t), t/(t + v)^2)$ with density

$$\hat{f}(m | x, t) = \frac{t + v}{\sqrt{2\pi t}} \exp\left(-\frac{t(m - x)^2}{2} - v(m - x)(m - \mu) - \frac{v^2(m - \mu)^2}{2t}\right).$$

When the receiver observes (m, x) , the inference on t is therefore a function of $(m - x)^2$ and $(m - \mu)^2$. In particular, \hat{f} satisfies the MLRP, so that a partial increase in $(m - x)^2$ is worse news about t , while a partial increase in $(m - \mu)^2$ is better news about t . It follows that $W(m | x)$ can be written as a function $\hat{W}(|m - x|, |m - \mu|)$, with a negative partial derivative in its first argument, and a positive partial derivative in its second argument. Since $V(m | s, t) = \int_X \hat{W}(|m - x|, |m - \mu|)q(x | s, t) dx$, we have

$$V_m(m | s, t) = \int_X \hat{W}_1(|m-x|, |m-\mu|) \operatorname{sgn}(m-x) q(x | s, t) dx + \operatorname{sgn}(m-\mu) \int_X \hat{W}_2(|m-x|, |m-\mu|) q(x | s, t) dx,$$

where $\operatorname{sgn}(z)$ denotes the sign of z , i.e., 1 when $z > 0$ and -1 when $z < 0$. Notice that $\hat{W}_1(|m-x|, |m-\mu|) \operatorname{sgn}(m-x)$ is an odd function of x around m . Since the normal posterior $q(x | s, t)$ is symmetric around $m(s, t)$, we can conclude that the first integral vanishes. On the other hand, $\hat{W}_2 > 0$, so the second line is proportional to $\operatorname{sgn}(m-\mu)$. This is positive if $m > \mu$, negative if $m < \mu$. *Q.E.D.*

Proof of Proposition 7. We establish this result through two lemmas. Lemma A2 shows with a simple fixed-point argument that there exists some interior iso-posterior curve with indifference $V(m | s, t) = V(m' | s, t)$ among the messages above and below the same curve. Lemma A3 shows, for incentive compatibility, that $V(m | s, t) - V(m' | s, t)$ is monotonic across iso-posterior curves. *Q.E.D.*

Lemma A2. There exists a $k \in (-1, 1)$ such that $V(m | s, t = k/s) = V(m' | s, t = k/s)$ when m is sent by $st \leq k$ types and m' is sent by the others.

Proof. First, we argue that for k sufficiently close to 1, message m' is better than m . As $k \rightarrow 1$, the set m' of (s, t) satisfying $st \geq k$ shrinks toward the corner $(s, t) = (1, 1)$. Since v is strictly increasing and $p(t)$ is nondegenerate, there exists some $t^* \in (0, 1)$ with $v(t^*) > Ev(t)$. When $k > t^*$, message m' is sent only by types who know $t > t^*$, so that $V(m' | s, t) > v(t^*)$. On the other hand, as k tends to 1, m tends toward an uninformative message, so $V(m | s, t) \rightarrow Ev(t)$. In conclusion, for large k , $V(m | s, t) < v(t^*) < V(m' | s, t)$. This is true for all pairs (s, t) , and in particular for all those with $st = k$.

By analogy, when k is sufficiently close to -1 , senders with $st = k$ prefer message m over m' . When k changes, we continuously change $V(m | s, t = k/s)$ and $V(m' | s, t = k/s)$. We know that $V(m | s, t = k/s) - V(m' | s, t = k/s)$ is positive for k near -1 and negative for k near $+1$. There must be an intermediate equilibrating k . *Q.E.D.*

Lemma A3. If the receiver believes that m is sent by the $st \leq k$ types and that m' is sent by the others, then $V(m' | s, t = \ell/s) - V(m | s, t = \ell/s)$ is increasing in ℓ .

Proof. We first prove that $P(t | m, x)$ increases in x when $k \geq 0$. From

$$\hat{f}(m | x, t) = \int_{-1}^{k/t} f(s | x, t) ds = \int_{-1}^{k/t} \frac{1+stx}{2} ds = \frac{I_1(t) + I_2(t)tx}{2},$$

where $I_1(t) = \int_{-1}^{k/t} ds = 2$ for $0 \leq t \leq k$ and $I_1(t) = 1 + k/t$ for $k \leq t \leq 1$ and, similarly, $I_2(t) = \int_{-1}^{k/t} s ds = 0$ for $0 \leq t \leq k$ and $I_2(t) = (k^2/t^2 - 1)/2$ for $k \leq t \leq 1$. We have $\hat{f}(m | x) = \int_0^1 \hat{f}(m | x, t) p(t) dt = (E[I_1(t)] + E[I_2(t)tx])/2$, so that

$$p(t | m, x) = \frac{I_1(t) + I_2(t)tx}{E[I_1(t)] + E[I_2(t)tx]} p(t)$$

and

$$P(t | m, x) = \frac{E[I_1(t) | t' \leq t] + E[I_2(t)tx | t' \leq t]x}{E[I_1(t)] + E[I_2(t)tx]} P(t). \quad (\text{A1})$$

Since I_2 is nonpositive and decreasing in t , we have $0 \geq E[I_2(t)tx | t' \leq t] \geq E[I_2(t)tx]$. Similarly, $I_1(t)$ is positive and decreasing in t so that $E[I_1(t) | t' \leq t] \geq E[I_1(t)] \geq 0$. According to (A1), $P(t | m, x)$ is then increasing in x . An analogous argument shows that $P(t | m, x)$ is again increasing in x when $k < 0$. Finally, since $v(t)$ is increasing in t , $W(m | x) = \int_T v(t) p(t | m', x) dt$ is decreasing in x . From (8), an increase in ℓ yields a first-order stochastic dominance increase in $q(x | s, t = \ell/s)$, so that in turn $V(m | s, t = \ell/s) = \int_X W(m | x) q(x | s, t = \ell/s) dx$ decreases in ℓ .

Similar calculations show $P(t | m', x)$ decreasing in x and $V(m' | s, t = \ell/s)$ increasing in ℓ . We conclude that $V(m' | s, t = \ell/s) - V(m | s, t = \ell/s)$ is increasing in ℓ . *Q.E.D.*

Proof of Proposition 8. The reputational value of message m^i for expert i with signal realization s^i is

$$V^i(m^i | s^i) = \int_X W^i(m^i | x) q(x | s^i) dx, \quad (\text{A2})$$

where

$$W^i(m^i | x) = \int_{M^{-i}} \int_{T^i} \int_{T^{-i}} u^i(t^i, t^{-i}) p(t^{-i} | m, x) dt^{-i} p(t^i | m, x) dt^i \hat{f}(m^{-i} | x) dm^{-i}.$$

By the assumption of independence of the ability of different experts, posterior reputations are again stochastically independent. In particular, the posterior reputation of an expert is stochastically independent of the message reported by another expert, $p(t^{-i} | m^i, m^{-i}, x) = p(t^{-i} | m^{-i}, x)$ and $p(t^i | m^i, m^{-i}, x) = p(t^i | m^i, x)$, so that

$$\begin{aligned} W^i(m^i | x) &= \int_{T^i} \int_{M^{-i}} \int_{T^{-i}} u^i(t^i, t^{-i}) p(t^{-i} | m^{-i}, x) dt^{-i} \hat{f}(m^{-i} | x) dm^{-i} p(t^i | m^i, x) dt^i \\ &= \int_{T^i} v^i(t^i, x) p(t^i | m^i, x) dt^i, \end{aligned}$$

where

$$v^i(t^i, x) = \int_{T^{-i}} u^i(t^i, t^{-i}) \int_{M^{-i}} p(t^{-i} | m^{-i}, x) \hat{f}(m^{-i} | x) dm^{-i} dt^{-i}.$$

The law of iterated expectations gives

$$\int_{M^{-i}} p(t^{-i} | m^{-i}, x) \hat{f}(m^{-i} | x) dm^{-i} = p(t^{-i} | x) = p(t^{-i}),$$

where we also used the independence of t^{-i} and x . It follows that

$$v^i(t^i, x) = \int_{T^{-i}} u^i(t^i, t^{-i}) p(t^{-i}) dt^{-i},$$

and thus $v^i(t^i, x)$ does not depend on x . Furthermore, since u^i is increasing in t^i for any t^{-i} , we find that $v^i(t^i)$ is an increasing function of t^i . We are thus back to the original problem with absolute reputational concerns, with the individual objective function $v^i(t^i) = E_{t^{-i}} u^i(t^i, t^{-i})$.

When expert i knows her own type, equation (A2) becomes

$$V^i(m^i | s^i, t^i) = \int_X W^i(m^i | x) q(x | s^i, t^i) dx$$

and the rest of the proof goes through as before. *Q.E.D.*

References

- AVERY, C.N. AND CHEVALIER, J.A. "Herding over the Career." *Economics Letters*, Vol. 63 (1999), pp. 327–333.
- BATTAGLINI, M. "Multiple Referrals and Multidimensional Cheap Talk." *Econometrica*, Vol. 70 (2002), pp. 1379–1401.
- BAYARRI, M.J. AND DEGROOT, M.H. "Gaining Weight: A Bayesian Approach." In J.M. Bernardo, M.H. DeGroot, D.V. Lindley, and A.F.M. Smith, eds., *Bayesian Statistics 3*. Oxford: University Press, 1988.
- AND ——— "Optimal Reporting of Predictions." *Journal of the American Statistical Association*, Vol. 84 (1989), pp. 214–222.
- BENABOU, R. AND LAROQUE, G. "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility." *Quarterly Journal of Economics*, Vol. 107 (1992), pp. 921–948.
- BERNHEIM, B.D. "A Theory of Conformity." *Journal of Political Economy*, Vol. 102 (1994), pp. 841–877.
- BRANDENBURGER, A. AND POLAK, B. "When Managers Cover Their Posteriors: Making the Decisions the Market Wants to See." *RAND Journal of Economics*, Vol. 27 (1996), pp. 523–541.
- CAL, H. AND WANG, J.T. "Overcommunication in Strategic Information Transmission Games." *Games and Economic Behavior*, forthcoming.
- CAMPBELL, C. "Learning and the Market for Information." Mimeo, Ohio State University, 1998.
- CRAWFORD, V. AND SOBEL, J. "Strategic Information Transmission." *Econometrica*, Vol. 50 (1982), pp. 1431–1452.
- DEWATRIPONT, M., JEWITT, I., AND TIROLE, J. "The Economics of Career Concerns, Part I: Comparing Information Structures." *Review of Economic Studies*, Vol. 66 (1999), pp. 183–198.
- EFFINGER, M.R. AND POLBORN, M.K. "Herding and Anti-Herding: A Model of Reputational Differentiation." *European Economic Review*, Vol. 45 (2001), pp. 385–403.
- GEANAKOPOLOS, J., PEARCE, D., AND STACCHETTI, E. "Psychological Games and Sequential Rationality." *Games and Economic Behavior*, Vol. 1 (1989), pp. 69–79.
- GENEST, C. AND ZIDEK, J.V. "Combining Probability Distributions: A Critique and an Annotated Bibliography." *Statistical Science*, Vol. 1 (1986), pp. 114–148.
- GREEN, J.R. AND STOKEY, N.L. "A Two-Person Game of Information Transmission." Discussion Paper no. 751, Harvard Institute of Economic Research, 1980.
- HOLMSTRÖM, B. "On the Theory of Delegation." In M. Boyer and R. Kihlstrom, eds., *Bayesian Models in Economic Theory*. Amsterdam: Elsevier, 1984.
- . "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies*, Vol. 66 (1999), pp. 169–182.
- AND RICART I COSTA, J. "Managerial Incentives and Capital Management." *Quarterly Journal of Economics*, Vol. 101 (1986), pp. 835–860.

- KRISHNA, V. AND MORGAN, J. "A Model of Expertise." *Quarterly Journal of Economics*, Vol. 116 (2000), pp. 747–775.
- LEHMANN, E.L. "Comparing Location Experiments." *Annals of Statistics*, Vol. 16 (1988), pp. 521–533.
- LEVY, G. "Anti-herding and Strategic Consultation." *European Economic Review*, Vol. 48 (2004), pp. 503–525.
- LI, W. "Mind Changes and the Design of Reporting Protocols." Mimeo, University of California at Riverside, 2004.
- OSBAND, K. "Optimal Forecasting Incentives." *Journal of Political Economy*, Vol. 97 (1989), pp. 1091–1112.
- OTTAVIANI, M. AND PRAT, A. "The Value of Public Information in Monopoly." *Econometrica*, Vol. 69 (2001), pp. 1673–1683.
- AND SØRENSEN, P. "Herd Behavior and Investment: Comment." *American Economic Review*, Vol. 90 (2000), pp. 695–704.
- AND ———. "Information Aggregation in Debate: Who Should Speak First?" *Journal of Public Economics*, Vol. 81 (2001), pp. 393–421.
- AND ———. "Professional Advice." *Journal of Economic Theory*, Vol. 126 (2006), pp. 120–142.
- AND ———. "The Strategy of Professional Forecasting." *Journal of Financial Economics*, forthcoming.
- PARTHASARATHY, K.R. *Probability Measures on Metric Spaces*. New York: Academic Press, 1967.
- PRAT, A. "The Wrong Kind of Transparency." *American Economic Review*, Vol. 95 (2005), pp. 862–877.
- PRENDERGAST, C. AND STOLE, L. "Impetuous Youngsters and Jaded Oldtimers: Acquiring a Reputation for Learning." *Journal of Political Economy*, Vol. 104 (1996), pp. 1105–1134.
- SCHARFSTEIN, D. AND STEIN, J. "Herd Behavior and Investment." *American Economic Review*, Vol. 80 (1990), pp. 465–479.
- SEIDMANN, D. "Effective Cheap Talk with Conflicting Interests." *Journal of Economic Theory*, Vol. 50 (1990), pp. 445–458.
- SOBEL, J. "A Theory of Credibility." *Review of Economic Studies*, Vol. 52 (1985), pp. 557–573.
- TRUEMAN, B. "Analyst Forecasts and Herding Behavior." *Review of Financial Studies*, Vol. 7 (1994), pp. 97–124.
- ZARNOWITZ, V. "An Appraisal of Short-Term Economic Forecasts." Mimeo, NBER, 1967.
- ZITZEWITZ, E. "Opinion-Producing Agents: Career Concerns and Exaggeration." Mimeo, Stanford Graduate School of Business, 2001.