Biased Advice in Dynamic Consulting*

Markus Reisinger[†] & Ramon Xifré[‡]

February 2023
Preliminary and Incomplete Version

Abstract

We study a two-period relationship between a client and a consultancy. In each period, the client's goal is to estimate the realization of a random variable but he is uncertain about his ability to do so. He can obtain costly advice from the consultancy. The report of the consultancy does not only influence the client's estimate about the variable but also provides information about his ability. We show that the consultancy biases her report away from the client's signal to lower the client's confidence in his ability and obtain a higher price for her advice in the future. In equilibrium, the client does not fully learn the consultancy's signal but for extreme cases. We also show that the relationship between the two parties may end after the first period if the consultancy does not obtain precise information about the client's signal. In this case, the consultancy is not fully informed about the client's self-confidence and demands a too high price for her advice.

Keywords: signaling, consultant-client relationship, advice, biased reporting.

JEL classification numbers: C72, D82.

*We wish to thank Marco Celentani, Luis Corchón, Riccardo Martina, Andreu Mas-Colell, Marc Möller, Pablo Ruiz-Verdú, Klaus Schmidt, Armin Schmutzler, and Mike Wiper for comments on previous drafts. The usual disclaimer applies.

[†]Frankfurt School of Finance & Management, Economics Department, Adickesallee 32-34, 60322 Frankfurt am Main, Germany. E-Mail: m.reisinger@fs.de.

[‡]ESCI-UPF School of International Studies and Public-Private Sector Research Center, IESE Business School, Pg. Pujades 1, 08003 Barcelona, Spain. E-Mail: ramon.xifre@esci.upf.edu

1 Introduction

In many industries, companies repeatedly use the advice of consultants. For example, strategic consultancies, such as Boston Consulting Group or McKinsey, advice businesses in different industries about decisions regarding the product portfolio, the targeted customer groups, the pursuing of a merger, etc. This is valuable to the management of the client company because it reduces uncertainty about the profitability of following a particular strategy or implementing a new project.

A similar relation holds between companies who gain access to a large set of consumer data through their business operations and data analytics companies. The former company often has a vague idea on how to use the data for advertising or pricing purposes but does not exactly know what is the best way to use the data and is uncertain about its own expertise in this domain. The latter company, due to its expertise in data technology, can give advice how to exploit the data in the most profitable way with respect to marketing strategies and pricing tactics. Indeed, companies such as Bloomberg or Axciom often use machine learning algorithms and other modern data analytic methods to provide consulting service. Such relationships become more and more prevalent as the possibility to collect data gets more abundant.

These consultant-client relationships are characterized by three important features: First, during the consulting process, the consultancy obtains information about the strategy that the client would have pursued without interaction with the consultancy. This could either be due to observing the current actions of the client or because providing the relevant information for successful consulting reveals this information fully or partly. Second, the relationship between the two companies is often a repeated one. The business environment will be subject to change in the future and/or consumer data will be updated and new possibilities for marketing and selling may come up. This implies that advice from the consultancy will be valuable again. Third, the consultancy has valuable information for undertaking the project and is usually more confident about its own ability than the client. The client's uncertainty is usually due to its lack of experience and expertise in the field (i.e., typically it is exploring unknown ground). By contrast, the consultancy is specialized in the task and, for example, thanks to its experience in consulting in similar projects has a more accurate estimation of its ability in dealing with the current project.

The aim of this paper is to provide a framework to study these issues and analyze the consequences for the truthfulness of the consultancy's advice and the efficiency of the relationship. We set up a two-period model in which a client (he) is uncertain about his ability to perform a given task (such as pursuing the best merger strategy or gaining most profits from consumer data). The company can hire a consultancy (she) to provide advice. The report send by the consultancy in the first period then has two effects: (i) it helps the client to make better choices about the current project and (ii) allows him to learn about his ability. The updated ability then determines the valuation for the client to work with the consultancy again in the second period. As this is anticipated in equilibrium, the consultancy may have the incentive to strategically misreport.

We think that our framework captures important traits of a variety of consultant-client relationships. The management literature has discussed extensively the role consultants play in selling security or providing reassurance to clients. It is a key rationale for understanding the whole industry from a psychodynamic view. Although this is a complex issue (Sturdy, 1997; Fincham, 1999; Mohe and Seidl, 2011), there is evidence that at least in some consultant-client relationships, ''both the adoption and discarding of ideas are based on largely subconscious processes—managerial anxiety over the uncertainty surrounding their careers, work role and organizational environment'' (Sturdy, 1997). With our model, the goal is to capture formally some of these effects.

We consider a model with two periods in which a client can hire a consulting company to provide advice about the state of the world (i.e., the profitability of a project, the optimal use of data, etc.). Both the client and the consultant receive a signal about the state. In addition, the consultancy obtains information about the client's signal as well. These assumptions follow Prendergast (1993). Differently to his analysis, in our model, the client—in contrast to the consultancy—is uncertain about the precision of his signal. This captures in a natural way the effect that the client is uncertain about his own ability. Therefore, after receiving the report of the consultancy, the client not only updates his estimation of the state of the world but also his ability to predict the state of the world correctly. A theoretical contribution of the paper is to provide a signal structure that allows updating of the precision of the client's signal in a tractable way.

In each period, the client and the consultancy bargain over the price for consulting services. After signals are drawn, the consultancy sends a message to the client. The consultancy can distort the message away from the signal she received but doing so is costly. This reflects the natural idea that while misreporting is possible, the consultancy needs to spend additional time and effort to credibly report that the state of the world is further away from the estimated value (i.e. misreporting that a project is profitable is more costly if it is expected to be highly unprofitable rather than if it expected to be only slightly unprofitable). The consultancy pursues two goals when sending the report. The first is the direct goal to reduce costs, which, *ceteris paribus*, induces the consultancy to report truthfully.[1] The second goal is to strategically influence the client's estimation about his ability. Specifically, a lower belief of the client's ability implies that the client (in expectation) benefits more from the consultancy's advice in the future and is therefore willing to pay a higher price. As a consequence, the consultancy has the incentive to distort her report to lower the client's belief about his ability.

In the second (last) period, the strategic effect is not present because there is no rehiring of the consultancy, which implies that there is truthful reporting. This does not hold in the first period. We show that there is a unique equilibrium in which the consultancy always misreports. The consultancy misreports by sending a message in the opposite direction of what she believes the client has observed.[2] She does so to undermine the client's belief in his ability, which allows the consultancy to obtain a higher revenue from the client in the next period. We therefore obtain the opposite result to Prendergast (1993)—i.e., the consultancy has the incentive to disagree with the client instead of confirming his estimate.

The resulting equilibrium is unique and fully separating, that is, each type (i.e., each consultancy with a different signal) sends a different report to the client. Since, in general, the consultancy gets only a noisy signal of the client's signal, the equilibrium is not fully revealing about what the consultant has observed on about the project. However, in the two extreme cases in which the consultancy learns the client's signal with certainty and when the consultancy does not learn the client's signal at all, the equilibrium is fully revealing. In these cases, although the consultancy distorts the signal, the client fully understands how to interpret the message and learns the truth. The distortion that occurs in equilibrium is therefore only costly to the consultancy but does not help her to increase the payment in next period.

[1]This could also capture the value of the consultancy's reputation in future projects vis-à-vis third parties.

[2]For example, if the consultancy receives the signal that the state of the world is 3 and her expectation about the client's signal is 1, she will send a report with a number larger than 3.

We show that the probability with which the consultancy and the client reach an agreement in the next period, depends on the noise of the signal that the consultancy obtains about the client's signal. Specifically, the lower the noise, the more precise the estimate that the consultancy has about the client's self-confidence, and therefore about his benefit from working with the consultancy again. If the noise is large, the probability that the consultancy claims a payment, which exceeds the client's benefit from consulting services, is relatively high. In this case relationship between the two parties ends after the first period.

Our paper is related to three main strands of the literature: subjective performance evaluation; dynamic principal-agent setups; and strategic reporting. We discuss the relations to these literatures in turn and also to previous papers studying the consulting business from a theoretical perspective.

First, a large literature studies optimal contracting under subjective performance evaluation (e.g., Prendergast 1993; Baker et al., 1994; MacLeod, 2003; Gibbs et al. 2004). This literature provides conditions under which efficient contracts are possible when evaluations are subjective.[3] As mentioned above, some elements of the basic information and reporting structure in the present paper are taken from Prendergast (1993). The main point of Prendergast (1993) is to show that subjective performance evaluation generates an agent's desire to conform to the principal's views.[4] His model was static and includes effort choice for the agent to acquire information.[5] The paper finds that the principal is facing a trade-off between inducing the agent to exert effort and inducing him to tell the truth. In contrast to this literature, we show that subjective performance evaluation in a dynamic set-up may distort the incentives to report truthfully, but in the opposite direction. Here we argue that in case there is uncertainty about the principal's ability, the agent has an incentive to contradict the principal's opinion and jeopardize her learning process.

Second, in most dynamic models, a party seeks to manipulate another party's beliefs because the first party's ability is uncertain and private information. For example, Prendergast and Stole (1996) show that when a manager's reputation is at

[3] Bellemare and Sebald (2018) test experimentally how over- and under-confident workers react to subjective performance measures.

[4] Klein and Mylovanov (2017) also show the tendency to conformism but in a dynamic setup.

[5] To simplify the analysis, in this paper we abstract from effort choices and we focus on the incentives for truthful reporting. In this respect, the relationships are similar to those between firms and leaders (Aghion and Jackson, 2016), in which wrong actions are rarely the result of poor effort but rather of inadequate information and career concerns, as it will be the case in the present paper.

stake, he may wish to appear as a fast learner, or Holmström (1999) demonstrates that when the agent's talent is not fully known, he will overinvest to acquire ability. The returns from building a reputation are highest when there is more uncertainty about the agent's type. Similarly, Levy (2005) finds that careerist judges contradict precedents too often in order to signal their ability. In our model, however, the motivation for the consultancy to report strategically comes from the uncertainty in the principal's ability or talent.[6]

Third, some papers that have studied the incentives for agents to misreport their information in setups without effort choices (Fudenberg and Tirole, 1986; Fischer and Verrecchia, 2000; Esö and Szentes, 2007; Iossa and Jullien, 2007). However, none of these papers considers the incentive to bias the report to influence the confidence of the other party in his own ability.[7]

The theoretical literature on client-consulting relationships is slim. Esö and Szentes (2007) develop a model in which the client may wish to keep the consultant (partially) in the dark so that she does fully understand the importance or the meaning of the task she has been commissioned. Bergemann et al. (2018) study the optimal selling mechanism of a data seller and show that it consists of a menu of statistical experiments. None of these papers considers the effect of information provision on the confidence of the client's ability.

The paper is organized as follows. Section 2 we introduce the model and notation. Section 3 analyzes the belief updating process. In Section 4, we present the main analysis and results. Finally, Section 5 concludes. All proofs are relegated to the Appendix.[8]

2 The model

Parties and Information: There are two firms, a consultancy company C and a buyer/client B, and two periods, $t = 1, 2$. The buyer seeks to obtain information about the true value of a parameter, η_t, in each period. In particular, η_t is a random

[6]Kaya (2010) considers a model in which the principal delays the acquisition of information about an agent's effort cost to incentivize the agent to exert effort. By contrast, in our paper the learning process is about the own ability.

[7]Ishida (2012) studies self-esteem in a one-period model. However, this is done by incorporating explicitly a component in the agent's utility function that captures his benefits from his ability type. Our model differs from this paper because self-confidence concerns appear endogenously.

[8]Because the paper is in a preliminary and incomplete stage, the current version only provides the proofs of Lemma 1 and 2. The proofs of the propositions are missing. They are not typed in yet but are already completed.

variable and we denote its realization by η_t^*. The random variables η_1 and η_2 are independent. In each period, the common prior is that η_t is distributed normally according to $\mathrm{N}(\eta_{0t}, \sigma_0^2)$ where η_{01}, η_{02} and σ_0^2 are common knowledge. Realizations η_1^* and η_2^* are only observed later in time.

In period t, both parties receive an imperfect, independent signal about η_t^* and C also receives a signal of B's signal.[9] More specifically, B observes η_{Bt}^*, which is the realization of $\eta_{Bt} = \eta_t^* + \varepsilon_B$, where ε_B is a random variable. Similarly, C observes η_{Ct}^*, which is the realization of $\eta_{Ct} = \eta_t^* + \varepsilon_C$, and also η_{ct}^* which is the realization of $\eta_{ct} = \eta_{Bt}^* + \varepsilon_c$, where ε_C and ε_c are two independent random variables. All signals η_{Bt}^*, η_{Ct}^*, and η_{ct}^* are non-verifiable—however, signal η_{Ct} may be revealed during the game (as will be explained below).

The assumption that the consulting firm has information on the client's opinion is realistic in most consultancy cases. For example, the data analytics company can observe the advertising or marketing strategy that the client is currently following, which represents the current state of information of the buyer. Similarly, consultancies know (at least imperfectly) how the firm they are consulting for plans to proceed without consulting.

We assume that $\varepsilon_C \sim \mathrm{N}(0, \sigma_C^2)$ and $\varepsilon_c \sim \mathrm{N}(0, \sigma_c^2)$ in both periods, with σ_C^2 and σ_c^2 given and common knowledge. We introduce uncertainty on the variance of the error term of the buyer's signal, $\sigma_B^2 = \mathrm{Var}(\varepsilon_B)$. Nature chooses this variance that can be low, σ_L^2, or high, σ_H^2, with $\sigma_H^2 > \sigma_L^2$, but the realization is not revealed to either the buyer or to the consulting firm. We write

$$\varepsilon_B \sim \theta \mathrm{N}(0, \sigma_L^2) + (1 - \theta)\mathrm{N}(0, \sigma_H^2), \tag{1}$$

where θ is a random variable with support [0,1] that can be interpreted as a measure of the buyer's self-confidence in his ability to interpret the data in the correct way. Variances σ_L^2 and σ_H^2 are given and common knowledge.

To allow for the computation of closed-form posterior distributions of θ, we further assume that θ has a prior distribution which is conjugate with the normal; that is, a distribution such that its posterior, if updated with normally distributed observations, remains in the same family of the prior distribution. In particular we assume that θ has *a priori* a Beta distribution with parameters α and β, $\mathrm{Be}(\alpha, \beta)$, and for simplicity α and β are restricted to be integer numbers. Any reasonably smooth unimodal distribution on [0,1] is likely to be well approximated by some

[9]The information structure in both periods is similar to the static one in Prendergast (1993).

Beta distribution. Specifically, if $\theta \sim$Be(α, β), then E$[\theta] = \alpha/(\alpha+\beta)$ and Var$(\theta) = \alpha\beta/((\alpha+\beta)^2(\alpha+\beta+1))$. For instance, the case of $\alpha = \beta = 1$ reproduces the uniform distribution on [0,1]; if $\alpha > \beta$ it is skewed to the right with E$[\theta] > 1/2$, and vice versa. On the other hand, the larger the sum of α and β the smaller the variance.[10]

Under this specification, the expected value of θ, E$[\theta]$, is a measure of the buyer's expected ability to deal with the information at each period. Notice that, using the formula for the conditional variance,

$$\begin{aligned} \text{Var}(\varepsilon_B) &= \text{Var}_\theta\left(\text{E}\left[\varepsilon_B|\theta\right]\right) + \text{E}_\theta\left[\text{Var}\left(\varepsilon_B|\theta\right)\right] \\ &= \sigma_H^2 - (\sigma_H^2 - \sigma_L^2)\text{E}[\theta], \end{aligned} \tag{2}$$

which makes it explicit that the variance of the error term of the buyer's signal is decreasing in his self-confidence.

Reporting: After observing her signal and, with some noise, the one of the client (i.e., η^*_{Ct} and η^*_{ct}), C sends a report, denoted by η_{Rt}, to B. The consultancy C can distort her report at a cost. This cost equals

$$k\left[\eta_{Rt} - \eta^\star_{Ct}\right]^2.$$

This falsification technology implies that the cost from misreporting is the larger the further away the report from the actual observation.

Negotiation: In each period t, B and C bargain over a payment p_t paid from B to C. We assume a random-proposer take-it-or-leave-it bargaining, in which B proposes with probability γ and C with probability $1-\gamma$.[11] The payment p_t is not conditional on the realization of η_t in each period, either because it is only observed later or it is difficult to contract on. This is realistic in our context as the success of the advice given by the data analytics or consultancy company is revealed only with a considerable time lag. In addition, the optimal strategy following from η_t can be difficult to describe and verify in a court of law.

Payoffs: C's payoff is the payment p_t minus the costs of distorting her report in each period. For simplicity, we assume that there is no discounting.[12]

[10]See, for example, Lee (1997) for further details.

[11]This is equivalent to Nash bargaining between B and C, where B has weight γ and C has weight $1-\gamma$.

[12]Allowing for discounting does not change the qualitative results. We abstract from it here to focus on strategic (and not the intertemporal) issues of the report.

As in Prendergast (1993), B's payoff π is the negative of the variance of the estimate η_t in each period, minus the price p_t paid to the consultancy, in case the two parties reach an agreement. Therefore, $\pi_t = -\text{Var}(\hat{\eta}_t) - p_t$, where $\hat{\eta}_t$ is the buyer's estimate of η_t. An interpretation, following Prendergast (1993), is that η_t is the profit of a project accruing to the buyer. As the buyers wants to prepare and undertake only profitable projects, it is reasonable to assume that his goal is to have information about η as accurate as possible. In case C is a data analytics company providing advice which marketing or consumer targeting strategy the client should allow, a lower variance of η_t implies more precise information on this strategy, leading to a higher profit for the client.

Timing: The timing is as follows.

$t = 0$ Nature chooses $\sigma_B^2 \in \{\sigma_L^2, \sigma_H^2\}$, η_1 and η_2.

$t = 1.1$ B observes $\eta_{B1}^\star$; thereafter, B and C bargain over p_1.

$t = 1.2$ C observes $\eta_{C1}^\star$ and $\eta_{c1}^\star$, and reports η_{R1} to B.

$t = 1.3$ B builds the estimate $\hat{\eta}_1$, updates θ to $\theta|\eta_{R1}$, and both parties receive their period-1 payoffs.

$t = 2.1$ B observes $\eta_{B2}^\star$; thereafter, B and C bargain over p_2.

$t = 2.2$ C observes $\eta_{C2}^\star$ and $\eta_{c2}^\star$, and reports η_{R2} to B.

$t = 2.3$ B builds the estimate $\hat{\eta}_2$ and both parties receive their period-2 payoffs.

We will use weak perfect Bayesian equilibrium as our equilibrium concept as we need do not to introduce specific restrictions off the equilibrium path.

3 Updating

The report η_{Rt} by C is an additional piece of information obtained by B and induces B to update the distribution of η_t. Therefore, it has a direct, positive impact on his payoff. However, in this section we focus on a different effect of C's report. We will study how B learns about his ability by updating the distribution of the error term of his signal ε_B after receiving C's reports η_{R1} and η_{R2} assuming these are truthful, i.e. $\eta_{Rt} = \eta_{Ct}^*$. This prepares for the next section, in which we deal with the case of strategic reporting by C, that is, when we consider the possibility that C distorts her report and sends $\eta_{Rt} \neq \eta^*_{Ct}$.

As $\eta^*_{Bt} = \eta^*_t + \varepsilon_B$ and $\eta^*_{Ct} = \eta^*_t + \varepsilon_C$, the client can extract information about ε_B by computing the measure $z^*_t = \eta^*_{Bt} - \eta_{Rt} = \eta^*_{Bt} - \eta^*_{Ct}$ of the *discrepancy* between the realization of his signal and the report of the consulting company. Substituting the definitions above of η^*_{Bt} and η^*_{Ct}, we obtain that z^*_t is a realization of the random variable $z_t = \varepsilon_B - \varepsilon_C$, where ε_B distributed as in (1) and $\varepsilon_C \sim \mathrm{N}(0, \sigma^2_C)$, with σ^2_C constant and known. The client, by computing z^*_t, uses C's report to obtain a noisy signal on the noise of his own signal. For convenience, let us define $Z^*_t = (z^*_t)^2$, which is independent of the sign of z^*_t and is our actual measure of discrepancy.

Recall from (1) that ε_B is distributed as a mixture of known, normal distributions with the random weight θ following a Beta distribution. This coupled with the fact that ε_C is distributed normally ensures that the posterior distribution of θ remains in the Beta family. This, in turn, delivers a closed-form posterior for the distribution of ε_B, which helps to compute B's confidence in his ability, E[$\theta|\cdot$]. The following lemma deals with the updated distributions of θ after B receives C's reports (provided that B learns C's true signal η^*_{Ct} through the report).

Lemma 1 *Let $\theta \sim Be(\alpha, \beta)$.*

1. *B's posterior distribution of θ after receiving C's first-period report η^*_{R1} is*

$$\theta|Z^*_1 \sim p\left(Z^*_1\right) Be(\alpha+1, \beta) + (1 - p\left(Z^*_1\right))Be(\alpha, \beta+1) \tag{3}$$

where

$$p\left(Z^*_1\right) = \frac{\alpha}{\alpha + \beta\left(\sqrt{\frac{\sigma^2_L+\sigma^2_C}{\sigma^2_H+\sigma^2_C}}\exp\left\{\frac{\sigma^2_H-\sigma^2_L}{2(\sigma^2_H+\sigma^2_C)(\sigma^2_L+\sigma^2_C)}Z^*_1\right\}\right)} \tag{4}$$

2. *B's posterior distribution on θ after receiving C's second-period report η^*_{R2} is*

$$\begin{aligned}\theta|Z^*_1, Z^*_2 \quad \sim \quad & q(Z^*_2, p\left(Z^*_1\right))Be(\alpha+2, \beta) + r(Z^*_2, p\left(Z^*_1\right))Be(\alpha, \beta+2) + \\ & (1 - q(Z^*_2, p\left(Z^*_1\right)) - r(Z^*_2, p\left(Z^*_1\right)))Be(\alpha+1, \beta+1)\end{aligned} \tag{5}$$

where $\partial q(Z^*_2, p\left(Z^*_1\right))/\partial Z^*_2 < 0$ and $\partial r(Z^*_2, p\left(Z^*_1\right))/\partial Z^*_2 > 0$.

The result implies that when Z^*_t is low (i.e., η^*_{Bt} is relatively close to η^*_{Ct}), the probability that B's distribution of the error term is the the one with low variance is relatively high. The intuition for the updating process can be understood in the easiest way by considering the extreme case of $\eta^*_{Bt} = \eta^*_{Ct}$. This implies that a realization of B's signal of unknown variance coincides with the realization of another independent signal (contingent on η_t) of known variance. Such a coincidence is more likely to occur if signals η_{Bt} and η_{Ct} are more strongly correlated with each other—a situation that occurs with a higher probability if the correlation between each and η_t is higher. Given that the correlation between η_{Ct} and η_t is fixed, this can only occur if the correlation between η_{Bt} and η_t is high, which corresponds to the case of a signal with low variance σ^2_L. Then, as the distance between η^*_{Bt} and η^*_{Ct} increases, the distribution of the client's observation η_{Bt} is more likely to be a disperse one. This increases the odds for σ^2_H.

The property of the inference about θ maps directly into the updated confidence of B about his ability.

Lemma 2 $\partial Var[\varepsilon_B|Z^*_1]/\partial Z^*_1 > 0$ *for any* Z^*_1 *and* $\partial Var[\varepsilon_B|Z^*_1, Z^*_2]/\partial Z^*_2 > 0$ *for any* Z^*_1 *and* Z^*_2*.*

The result follows from the fact that when η^*_{Bt} and η^*_{Ct} are far from each other, the posterior distribution on θ places a low weight on the Beta distribution with the highest expected value. A large Z^*_t, therefore, decreases the weighted mean of the posterior distribution on θ and, therefore, B's confidence in his ability because it increases the conditional variance of ε_B. This finding is intuitive as it recommends B to revise downwards his confidence in the estimate of his own ability in case of observing a higher discrepancy between his private observation and another independent signal.

4 Analysis

After the explanation of the updating process, we now solve for the equilibrium of the game. We give a particular focus on the reporting strategy of C and whether or not B obtains perfect information about C's signal. We solve the model via backward induction (i.e., we start in the second stage).

4.1 Second Stage

As the consultant has costs from not reporting the truth in stage $t = 2.2$, but cannot gain from lying as the payment p_2 is already determined, she will report truthfully.

Therefore, if B and C reached an agreement in their negotiation on $t = 2.1$, B knows that he will receive a truthful report from C. This implies that the variance of his estimate $\hat{\eta}_{B2}$ is

$$\text{Var}(\hat{\eta}_{B2}) = \frac{\sigma_0^2 \sigma_{B2}^2 \sigma_C^2}{\sigma_0^2 \sigma_{B2}^2 + \sigma_0^2 \sigma_C^2 + \sigma_{B2}^2 \sigma_C^2}, \tag{6}$$

with $\sigma_{B2}^2 = \sigma_H^2 - (\sigma_H^2 - \sigma_L^2) E_2[\theta|\eta_{R1}]$, where, following the notation from the last section, $E_2[\cdot]$ is the principal's posterior about θ. Specifically, if B can deduce from the report η_{R1} that $\eta_{C1}^\star$ is close to $\eta_{B1}^\star$, he will update his posterior about θ upwards, which implies that σ_{B2}^2 is lower. This follows from Lemma 2.

We now determine the expected payment that the consulting firm obtains in the negotiation with the client. If B makes the offer (which occurs with probability γ), he will offer $p_2 = 0$, as C reports truthfully and therefore has no cost of misreporting. By contrast, with probability $1-\gamma$, C makes the offer and will demand the maximum payment from B. That is, a payment that leaves B indifferent between accepting the offer (and obtaining the report from C) and rejecting the offer (and receiving no report). When not reaching an agreement with C, B's variance of his estimate $\hat{\eta}_{B2}$ is

$$\text{Var}(\hat{\eta}_{B2}) = \frac{\sigma_0^2 \sigma_{B2}^2}{\sigma_0^2 + \sigma_{B2}^2}. \tag{7}$$

Because B's payoff is the negative of the variance, his gain from reaching an agreement with the agent is $(7) - (6)$. As a consequence, if C makes the offer, she can claim up to $(7) - (6)$ as payment. In equilibrium, $p_2 = (7) - (6)$, as B accepts when being indifferent.

Simplifying $(7) - (6)$, we obtain that C's expected payment in stage 2 is

$$(1-\gamma) \frac{\left(\sigma_0^2\right)^2 \left(\sigma_{B2}^2\right)^2}{\left(\sigma_0^2 \sigma_{B2}^2 + \sigma_0^2 \sigma_C^2 + \sigma_{B2}^2 \sigma_C^2\right) \left(\sigma_0^2 + \sigma_{B2}^2\right)}. \tag{8}$$

Taking the derivative with respect to σ_{B2}^2 yields

$$(1-\gamma) \frac{\left(\sigma_0^2\right)^3 \sigma_{B2}^2 \left(\sigma_0^2 \sigma_{B2}^2 + 2\sigma_C^2 (\sigma_0^2 + \sigma_{B2}^2)\right)}{\left(\sigma_0^2 \sigma_{B2}^2 + \sigma_0^2 \sigma_C^2 + \sigma_{B2}^2 \sigma_C^2\right)^2 \left(\sigma_0^2 + \sigma_{B2}^2\right)^2} > 0. \tag{9}$$

As a consequence C's payoff is strictly increasing in σ_{B2}^2. This implies that, ignoring the costs of misreporting, C has the incentive to set η_{R1} in the first stage so as to minimize $E_2[\theta|\eta_{R1}]$.

4.2 First Stage

In this section, we solve for the equilibrium of the full game. Before doing so, we consider in some detail the important benchmark in which C observes B signal with certainty, that is $\sigma_c^2 = 0$. In other words, C knows B's observation about the state of the world, η_1, which is given by $\eta_{B1}^\star$. This case not only serves as a benchmark but also provides insights into the working of the model, which can be used later on.

We first determine the type of equilibrium that can arise in this situation. We obtain the following result:

Proposition 1 *Any weak Perfect Bayesian Equilibrium of the game is a fully separating equilibrium. In this equilibrium, $\eta_{R1} \neq \eta_{C1}^\star$ for all types $\eta_{C1}^\star$.*

The first statement of Proposition 1 implies that there is no pooling (or partial pooling) equilibrium. The intuition behind this result is rooted in the continuum of types of C and that these types are unbounded. If a pooling equilibrium in which all types send the same report, say, $\bar{\eta}_{R1}$, existed, then each type with a signal realization $\eta_{C1}^\star$ sufficiently far away from $\bar{\eta}_{R1}$ would find it profitable to send a report closer to her signal than to send $\bar{\eta}_{R1}$. The reason is that, whatever B's belief is after the deviation, the fact that C has lower cost than from sending the report $\eta_{R1} = \bar{\eta}_{R1}$ dominates (due to the convexity cost of misreporting). A similar argument applies to any partial pooling equilibrium. As a consequence, the equilibrium must be a separating one.

From the second statement of Proposition 1, it follows that in the resulting separating equilibrium, no type will report truthfully. To understand this result, suppose to the contrary that a truthful reporting equilibrium would exist and consider C's incentive to deviate. Given that the expected payment in the second stage is given by (8), the parts of C's profit function, which depend on η_{R1} are

$$(1-\gamma)\frac{\left(\sigma_0^2\right)^2\left(\sigma_{B2}^2\right)^2}{\left(\sigma_0^2\sigma_{B2}^2+\sigma_0^2\sigma_C^2+\sigma_{B2}^2\sigma_C^2\right)\left(\sigma_0^2+\sigma_{B2}^2\right)} - k\left[\eta_{R1}-\eta_{C1}^\star\right]^2, \tag{10}$$

where $\sigma_{B2}^2 = \sigma_H^2 - (\sigma_H^2 - \sigma_L^2)\mathrm{E}[\theta|\eta_{R1}]$, with, using Lemma 1,

$$\mathrm{E}[\theta|\eta_{R1}] = \frac{\alpha}{\alpha+\beta+1}\left(1+\frac{1}{\alpha+\beta\left(\sqrt{\frac{\sigma_L^2+\sigma_C^2}{\sigma_H^2+\sigma_C^2}}\exp\left\{\frac{\sigma_H^2-\sigma_L^2}{2(\sigma_H^2+\sigma_C^2)(\sigma_L^2+\sigma_C^2)}\left(\eta_{R1}-\eta_{B1}^\star\right)^2\right\}\right)}\right). \tag{11}$$

This formula shows the main trade-off for C when deciding about her report. First, as is evident from the second term of (10), misreporting involves costs, which provides an incentive for C to report truthfully. At the same time, her report influences the payment she receives in the second period. Specifically, as pointed out above, we know from (9) that the first term of (10) is increasing in σ^2_{B2}, which in turn is increasing in the difference between η_{R1} and $\eta^\star_{B1}$. The latter follows because σ^2_{B2} falls in the expected value of θ and this expected value falls in $|\eta_{R1} - \eta^\star_{B1}|$. Therefore, sending a report which is further away from B's first-stage signal, increases C's expected second-stage payment.

In a truthful-reporting equilibrium, $\eta_{R1} = \eta^\star_{C1}$, which implies that the costs from misreporting are zero, a slight deviation from reporting $\eta_{R1} = \eta^\star_{C1}$ has negligible costs. By contrast, the gain is non-negligible as the derivative of the first term with respect to η_{R1} is non-zero. As a consequence, each type has an incentive to deviate in a truthful-reporting equilibrium. Therefore, in any separating equilibrium, $\eta_{R1} \neq \eta^\star_{C1}$.

We can now analyze the separating equilibrium in further detail. The first important observation is that, although the equilibrium does involve misreporting, the client obtains full information about the consultancy's signal $\eta^\star_{C1}$. The reason is that in case $\sigma^2_c = 0$, the only dimension of uncertainty for B is $\eta^\star_{C1}$. Because all types send a different report in the separating equilibrium, the client can back out the consultancy's true signal from the report. Therefore, in equilibrium, all information is transmitted, which implies that the equilibrium is efficient in this respect.

We now determine the direction of C's reporting bias. Doing so allows us to derive our next proposition.

Proposition 2 *In any equilibrium,* $|\eta_{R1} - \eta^\star_{B1}| > |\eta_{R1} - \eta^\star_{C1}|$.

Proposition 2 states that, in equilibrium, C distorts her report in a way that increases the difference between the report and B's signal $\eta^\star_{B1}$ compared to the difference between the report and C's signal $\eta^\star_{C1}$. In other words, for $\eta^\star_{C1} > \eta^\star_{B1}$, the client exaggerates the state of the world in her report whereas for $\eta^\star_{C1} < \eta^\star_{B1}$ she downplays it. The reason for this result can be easily seen in (10). Biasing the report in this direction increases the expected payment of C in the second stage in the least costly way and is therefore profitable. The intuition is that the consultancy has the incentive to lower the self-confidence of the client's ability. Her only way to do so is to send a report which is far away from the client's signal. Although the client can back out the true state of the world, incentive-compatibility of the consultancy's report can only be assured via such distortions as otherwise mimicking would be too cheap.

Proposition 2 therefore establishes a non-conformism result: The consultancy distorts her report further away from the client's signal than when reporting truthfully. This result is opposite to the one of Prendergast (1993). He obtains that if C (in his paper, the agent) is perfectly informed about B's (in his paper, the principal's) realization of the signal, C just reports B's realization $\eta^{\star}_{B1}$. This is 'Yes Man' in its strongest form. In addition, in this case B does not learn anything, which implies that no information is transferred from C to B. By contrast, in our model, B obtains full information, although all types distort their report away from B's signal.

We can also determine the extent of the distortion in a more precise way. Because in equilibrium C obtains full information, that is, the reports of the different types allow the client to infer the type of the consulting firm with certainty, we can use the profit function given by (10) to obtain how the report changes with the type of C. Following the same steps as in Mailath (1987) and Mailath and von Thadden (2013), we can show that the equilibrium strategy is differentiable, which implies that we can determine $d\eta_{R1}/d\eta^{\star}_{C1}$. Using the first-order condition of C's profit function, and defining

$$\Sigma\left(\sigma_0^2, \sigma_C^2, \sigma_{B2}^2, \sigma_H^2, \sigma_L^2, (\eta_{R1}-\eta^{\star}_{B1})^2\right) \equiv$$

$$\frac{\left(\sigma_0^2\right)^3 \sigma_{B2}^2 \left(\sigma_0^2\sigma_{B2}^2 + 2\sigma_0^2\sigma_C^2 + 2\sigma_{B2}^2\sigma_C^2\right)\left(\sigma_H^2-\sigma_L^2\right)^2 \sqrt{\frac{\sigma_L^2+\sigma_C^2}{\sigma_H^2+\sigma_C^2}} \exp\left\{\frac{\sigma_H^2-\sigma_L^2}{2(\sigma_H^2+\sigma_C^2)(\sigma_L^2+\sigma_C^2)}\left(\eta_{R1}-\eta^{\star}_{B1}\right)^2\right\}}{\left(\sigma_H^2+\sigma_C^2\right)\left(\sigma_L^2+\sigma_C^2\right)\left(\sigma_0^2+\sigma_{B2}^2\right)\left(\sigma_0^2\sigma_{B2}^2+\sigma_0^2\sigma_C^2+\sigma_{B2}^2\sigma_C^2\right)} > 0,$$

we can derive the following result:

Proposition 3 *There is a unique separating equilibrium, in which the optimal strategy is given by the differential equation:*

$$\frac{d\eta_{R1}}{d\eta^{\star}_{C1}} = \frac{\beta\left(\eta_{R1}-\eta^{\star}_{B1}\right)\left(E[\theta|\eta_{R1}]\right)^2 \Sigma\left(\sigma_0^2, \sigma_C^2, \sigma_{B2}^2, \sigma_H^2, \sigma_L^2, (\eta_{R1}-\eta^{\star}_{B1})^2\right)}{\alpha k\left|\eta_{R1}-\eta^{\star}_{C1}\right|} \tag{12}$$

where $E[\theta|\eta_{R1}]$ *is given by* (11). *It follows that* $d\eta_{R1}/d\eta^{\star}_{C1} > 0 (< 0)$ *if and only if* $\eta_{R1} > (<) \eta^{\star}_{B1}$.

Proposition 3 first establishes the intuitive result that the consultancy's report is increasing in her type if the type is above the client's signal and decreasing in the type if the type is below the client's signal. This can be seen in (12) because the sign of $d\eta_{R1}/d\eta^{\star}_{C1}$ depends on the sign of $\eta_{R1} - \eta^{\star}_{B1}$. Therefore, each type distorts her report in the opposite direction of the client's signal to lower the client's self-confidence.

The formula in Proposition 3 also allows us to determine how the extent of misreporting depends on the parameters. As an example, consider a change in σ_0^2. We obtain that the distortion of the report measured by $|d\eta_{R1}/d\eta^{\star}_{C1}|$ is stronger if σ_0^2 is larger. This is intuitive: If the prior distribution of the signal is more uncertain, the client has more difficulties interpreting his own signal with respect to his self-confidence. This implies that a consultancy's action of lowering the client's self-confidence has a stronger effect. As a consequence, the incentive to do so is larger, resulting in an increased distortion in equilibrium.

We finally note that, although information is fully transmitted, the equilibrium is inefficient. The reason is that the equilibrium involves misreporting, which is costly for the consultancy. The same information can be transmitted through truth-telling, and this outcome is less costly for the consultancy.

Having solved the benchmark case in which $\sigma_c^2 = 0$, we now turn to the general model in which $\sigma_c^2 > 0$, that is, C does not observe B's signal with certainty but only with some noise. We can then pursue a similar analysis as in the benchmark case to show that only a separating equilibrium exists, and that C's report differs from her observed signal. In contrast to the benchmark case, C needs to form an expectation about B's signal, which is given by

$$E\left[\eta^{\star}_{B1}\right] = \frac{\sigma_0^2\sigma_c^2\eta^{\star}_{C1} + \sigma_C^2\sigma_c^2\eta_0 + \sigma_0^2\sigma_C^2\eta^{\star}_{c1}}{\sigma_0^2\sigma_c^2 + \sigma_C^2\sigma_c^2 + \sigma_0^2\sigma_C^2}.$$

Given this expectation and the costs of distorting the report, the consultancy faces a similar trade-off as in the benchmark. Specifically, given $\eta^{\star}_{C1}$, C distorts the report in the opposite direction to her *expectation* of B's signal.

A main difference to the benchmark situation is that the client does not fully learn the consultancy's signal in equilibrium. Although the equilibrium is separating, B does not know $\eta^{\star}_{c1}$. He therefore faces a two-dimensional uncertainty about C's signals as neither $\eta^{\star}_{C1}$ nor $\eta^{\star}_{c1}$ is known. This prevents B from backing out the true signal of C in the separating equilibrium. For example, in case $\eta^{\star}_{R1} > \eta^{\star}_{B1}$—that is, C sends a high report—B does not know whether this is caused by a large signal of the consultancy (i.e., $\eta^{\star}_{C1} > \eta^{\star}_{B1}$) or by the belief of C that B's signal is relatively low (i.e., $\eta^{\star}_{c1} < \eta^{\star}_{B1}$). This result is summarized by the following proposition:

Proposition 4 *If $0 < \sigma^2_c < \infty$, there is a unique separating equilibrium in which B does not learn C's signal in the first period with certainty.*

As indicated by Proposition 4, the result that full learning does not occur in equilibrium depends on σ_c being finite. Instead, if $\sigma_c \to \infty$, the result is no longer true. If $\sigma_c \to \infty$, C does not learn anything about B's signal $\eta^{\star}_{B1}$. This implies that the expectation of type $\eta^{\star}_{C1}$ about $\eta^{\star}_{B1}$ is

$$\frac{\sigma^2_0 \eta^{\star}_{C1} + \sigma^2_C \eta_0}{\sigma^2_0 + \sigma^2_C}.$$

Similar to Prendergast (1993), the last expression has only one unknown from B's perspective (i.e., $\eta^{\star}_{C1}$), which implies that B can back out this information in the separating equilibrium. It follows that full learning of C's signal occurs only in the two extreme cases $\sigma^2_c = 0$ and $\sigma^2_c \to \infty$.

A major difference between the two extreme scenarios is that C is fully informed about B's self-confidence in case $\sigma^2_c = 0$ while this is not true if $\sigma^2_c \to \infty$. This implies that in the negotiation in stage 2, C cannot fully extract B's surplus in case it makes the offer. In particular, because C does not know $\eta^{\star}_{B1}$ with certainty, it needs to form an expectation about the self-confidence of B in the second period, and therefore also about the benefit that B receives from consulting services in the second period. Specifically, for any report η_{R1} send by C in the first period, the expected benefit for advice in the second period from C's perspective is

$$\mathrm{E}\left[\frac{\left(\sigma^2_0\right)^2 \left(\sigma^2_{B2}\right)^2}{\left(\sigma^2_0\sigma^2_{B2} + \sigma^2_0\sigma^2_C + \sigma^2_{B2}\sigma^2_C\right)\left(\sigma^2_0 + \sigma^2_{B2}\right)}\right] \tag{13}$$

with

$$\sigma^2_{B2} = \sigma^2_H - \frac{(\sigma^2_H - \sigma^2_L)\alpha}{\alpha+\beta+1}\left(1 + \frac{1}{\alpha + \beta\left(\sqrt{\frac{\sigma^2_L+\sigma^2_C}{\sigma^2_H+\sigma^2_C}}\exp\left\{\frac{\sigma^2_H-\sigma^2_L}{2(\sigma^2_H+\sigma^2_C)(\sigma^2_L+\sigma^2_C)}\left(\eta_{R1} - \eta^{\star}_{B1}\right)^2\right\}\right)}\right),$$

where the expectation is formed over $\eta^{\star}_{B1}$.

When making in offer in the second period, C's maximization problem is therefore given by

$$\max_{p_2} \; p_2 \, \text{prob} \left(p_2 \leq \frac{\left(\sigma_0^2\right)^2 \left(\sigma_{B2}^2\right)^2}{\left(\sigma_0^2 \sigma_{B2}^2 + \sigma_0^2 \sigma_C^2 + \sigma_{B2}^2 \sigma_C^2\right)\left(\sigma_0^2 + \sigma_{B2}^2\right)} \right).$$

As a consequence, due to uncertainty about B's benefit, C faces the well-known trade-off of demanding a larger payment and risking that B rejects the offer or asking for a low payment and getting the offer accepted with a higher probability. Since $\eta_{B1}^{\star}$ is not know by C with certainty, the probability that B rejects the offer is strictly positive. Therefore, in contrast to the situation in which C is fully informed about $\eta_{B1}^{\star}$, the two parties do not always reach an agreement in the second period when C cannot observe $\eta_{B1}^{\star}$ with certainty.

The next proposition summarizes the results explained above:

Proposition 5 *In the two extreme cases $\sigma_c^2 = 0$ and $\sigma_c^2 \to \infty$, B's signal in the first period is fully revealed.*
If $\sigma_c^2 = 0$, B and C reach an agreement with probability 1 in the second period, which implies that B's second-period signal is fully revealed. Instead, if $\sigma_c^2 \to \infty$, B and C reach an agreement in the second period with a probability strictly less than 1.

An immediate consequence of the result that the relationship between the client and the consultancy is terminated after one period with some probability is that in the second period valuable advice might get lost. Because the consultancy does not distort her advice in the second period, the consultancy's signal would be fully revealed to the client. This implies an additional inefficiency in the situation with $\sigma_c^2 \to \infty$ compared to the situation with $\sigma_c^2 = 0$. In both cases, an inefficiency in the first period occurs due to the distortion costs for the consultancy, and this inefficiency might be of different size. However, as the parties always reach an agreement in the second period in the situation with $\sigma_c^2 = 0$, there is full efficiency in the second period in this case but not with $\sigma_c^2 \to \infty$.

5 Conclusion

This paper studies a consultant-client relationship in a two-period framework. The client is uncertain about his ability to figure out the best strategy for his company or how to interpret data, and therefore lacks self-confidence. The consultancy, when sending a report to the client, does therefore not only induce the client to update his belief about the best strategy but also his belief about his ability to find the best strategy in the next period. This implies that the consultancy has the incentive to strategically report in such a way that the self-confidence of the client is undermined because it allows the consultancy to claim a higher payment for her services in the next period.

We show that there exists a unique equilibrium in the signaling game, which is fully separating. In this equilibrium, the consultancy distorts her report in the opposite direction of the (expected) signal of the client. This implies that the consultancy prefers to show non-conformism with the client because doing so may lowers the client's self-confidence. Due to this distortion, the information about the consultancy's signal is only imperfectly revealed to the client, apart from extreme cases. As the consultancy does not learn the self-confidence of the client with certainty, in the second period, she may demand a payment, which exceeds the client's valuation of the consultancy's advice. This leads to an additional source of inefficiency as valuable information may gets lost in this case.

There are several interesting directions that can be explored within the framework of our paper. First, it is possible to analyze how much information about his own signal the client is willing to reveal to the consultancy. This implies that one might think of σ_c^2 being endogenous and chosen by the client. The trade-off for the client is here that providing more certainty to the consultancy allows for better information transmission but also lowers his payoff in the second period. Second, the consultancy has an interest to learn the client's signal with a higher precision to obtain a larger payment in the second period. Therefore, she may spend effort to reduce σ_c^2. As a consequence, an endogenous σ_c^2 might be due to the interplay of the client's effort to hide the realization of his signal and consultancy's effort to obtain this information. Third, the extent of the distortion of the consultancy's report can be explored further as it is an important determinant in the efficiency of the relationship. Specifically, the effect of σ_c^2 on the distortion could be interesting to explore.

6 Appendix

Proof of Lemma 1

Part 1

We have that $\varepsilon_B \sim \theta\mathrm{N}(0,\sigma_L^2)+(1-\theta)\mathrm{N}(0,\sigma_H^2)$. The priors on θ are that $\theta \sim$Be(α,β). We are interested in the posterior distribution $\theta|Z_1$ where $Z_1 = (z_1)^2$ and $z_1 \sim N(0,\sigma_L^2+\sigma_C^2)$ with probability θ and $z_1 \sim N(0,\sigma_H^2+\sigma_C^2)$ with probability $1-\theta$.

First, notice that by Bayes' rule we have $f(\theta|Z_1) \propto f(Z_1|\theta)f(\theta)$. Therefore,

$$f(\theta|Z_1) \propto \theta^{\alpha-1}(1-\theta)^{\beta-1}\left(\theta\frac{\exp\left(-Z_1/(2(\sigma_L^2+\sigma_C^2))\right)}{\sqrt{\sigma_L^2+\sigma_C^2}}+(1-\theta)\frac{\exp\left(-Z_1/(2(\sigma_H^2+\sigma_C^2))\right)}{\sqrt{\sigma_H^2+\sigma_C^2}}\right).$$

Let $C_i(Z_1)=\exp\left(-Z_1/(2(\sigma_i^2+\sigma_C^2))\right)/\sqrt{\sigma_i^2+\sigma_C^2}$ for $i=L,H$. This leaves the above expression as

$$f(\theta|Z_1) \propto \theta^{\alpha-1+1}(1-\theta)^{\beta-1}C_L(Z_1)+\theta^{\alpha-1}(1-\theta)^{\beta-1+1}C_H(Z_1).$$

We now multiply and divide each term by a respectively suitable Beta function and we get

$$\begin{aligned} f(\theta|Z_1) \;&\propto\; \mathrm{B}(\alpha+1,\beta)C_L(Z_1)\frac{1}{\mathrm{B}(\alpha+1,\beta)}\theta^{\alpha-1+1}(1-\theta)^{\beta-1}+ \\ &\mathrm{B}(\alpha,\beta+1)C_H(Z_1)\frac{1}{\mathrm{B}(\alpha,\beta+1)}\theta^{\alpha-1}(1-\theta)^{\beta-1+1}. \end{aligned}$$

In each term one identifies a Beta distribution. So we are left with

$$\theta|Z_1 \sim p(Z_1)\mathrm{Be}(\alpha+1,\beta)+(1-p(Z_1))\mathrm{Be}(\alpha,\beta+1),$$

where

$$p(Z_1)=\frac{C_L(Z_1)\mathrm{B}(\alpha+1,\beta)}{C_L(Z_1)\mathrm{B}(\alpha+1,\beta)+C_H(Z_1)\mathrm{B}(\alpha,\beta+1)}.$$

Now, using the fact that $\mathrm{B}(\alpha,\beta)=\Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha+\beta)$ and that for α,β integers we have $\Gamma(n)=(n-1)!$, the expression for p above simplifies into

$$p(Z_1)=\frac{\alpha C_L(Z_1)}{\alpha C_L(Z_1)+\beta C_H(Z_1)}=\frac{\alpha}{\alpha+\beta\delta(Z_1)}, \tag{14}$$

where $\delta(Z_1)$, substituting back the expressions for $C_i(Z_1)$, is the following strictly increasing function of Z_1

$$\delta(Z_1) = \frac{C_H(Z_1)}{C_L(Z_1)} = \sqrt{\frac{\sigma_L^2 + \sigma_C^2}{\sigma_H^2 + \sigma_C^2}} \exp\left\{ \frac{\sigma_H^2 - \sigma_L^2}{2(\sigma_H^2 + \sigma_C^2)(\sigma_L^2 + \sigma_C^2)} Z_1 \right\}. \tag{15}$$

Part 2

We proceed in an analogous way to find the posterior $\theta|Z_2, Z_1$ when the priors are now that $\theta|Z_1 \sim p(Z_1)\text{Be}(\alpha+1,\beta) + (1-p(Z_2))\text{Be}(\alpha,\beta+1)$ for some Z_1 given. We have that $\theta|Z_2, Z_1 \propto f(Z_2|\theta)f(\theta|Z_1)$. We keep the definition of $C_i(Z_2)$ for $i = L, H$ from the previous part. This being the case, we can write

$$\begin{aligned} f(\theta|Z_2, Z_1) \;\propto\; & \left(p(Z_1)\frac{1}{\text{B}(\alpha+1,\beta)}\theta^{\alpha+1-1}(1-\theta)^{\beta-1} + \right. \\ & \left. (1-p(Z_1))\frac{1}{\text{B}(\alpha,\beta+1)}\theta^{\alpha-1}(1-\theta)^{\beta+1-1} \right) \times \\ & \left(\theta C_L(Z_2) + (1-\theta)C_L(Z_2)\right). \end{aligned}$$

Developing the product and multiplying and dividing each term by the appropriate Beta function, we get

$$\begin{aligned} f(\theta'|Z_2) \;\propto\; & p(Z_1)\frac{\text{B}(\alpha+2,\beta)}{\text{B}(\alpha+1,\beta)}\left(\frac{C_L(Z_2)}{\text{B}(\alpha+2,\beta)}\theta^{\alpha+2-1}(1-\theta)^{\beta-1}\right) + \\ & p(Z_1)\frac{\text{B}(\alpha+1,\beta+1)}{\text{B}(\alpha+1,\beta)}\left(\frac{C_H(Z_2)}{\text{B}(\alpha+1,\beta+1)}\theta^{\alpha+1-1}(1-\theta)^{\beta+1-1}\right) + \\ & (1-p(Z_1))\,\frac{\text{B}(\alpha+1,\beta+1)}{\text{B}(\alpha,\beta+1)}\left(\frac{C_L(Z_2)}{\text{B}(\alpha+1,\beta+1)}\theta^{\alpha+1-1}(1-\theta)^{\beta+1-1}\right) + \\ & (1-p(Z_1))\,\frac{\text{B}(\alpha,\beta+2)}{\text{B}(\alpha,\beta+1)}\left(\frac{C_H(Z_2)}{\text{B}(\alpha+1,\beta+1)}\theta^{\alpha-1}(1-\theta)^{\beta+2-1}\right). \end{aligned}$$

In the same vein as in part 1, we identify the corresponding Beta distributions and find that

$$\begin{aligned} f(\theta|Z_2, Z_2) \;\propto\; & \text{Be}(\alpha+2,\beta)\alpha p(Z_1)C_L(Z_2) + \text{Be}(\alpha+1,\beta+1)\beta p(Z_1)C_L(Z_2) + \\ & \text{Be}(\alpha+1,\beta+1)\alpha(1-p(Z_1))C_H(Z_2) + \text{Be}(\alpha,\beta+2)\beta(1-p)C_H(Z_2). \end{aligned}$$

It is convenient now to define,

$$\begin{aligned} q(Z_2, Z_1) &= \frac{\alpha p(Z_1) C_L(Z_2)}{(\alpha+\beta)\left(p(Z_1) C_L(z_2) + (1-p(Z_1)) C_H(z_2)\right)} \\ &= \frac{\alpha p(Z_1)}{(\alpha+\beta)\left(p(Z_1) + (1-p(Z_1))\delta(Z_2)\right)} \\ r(Z_2, Z_1) &= \frac{\beta(1-p(Z_1)) C_H(z_2)}{(\alpha+\beta)\left(p(Z_1) C_L(Z_2) + (1-p(Z_1)) C_H(Z_2)\right)} \\ &= \frac{\beta(1-p(Z_1))}{(\alpha+\beta)\left(p(Z_1)/\delta(Z_2) + 1 - p(Z_1)\right)} \end{aligned}$$

where $\delta(Z_2)$ is defined as in (15). Then we have that

$$\begin{aligned} \theta | Z_2, Z_1 &= q(Z_2, Z_1)\text{Be}(\alpha+2, \beta) + r(Z_2, Z_1)\text{Be}(\alpha, \beta+2) + \\ &\quad (1 - q(Z_2, Z_1) - r(Z_2, Z_1))\text{Be}(\alpha+1, \beta+1). \ \square \end{aligned}$$

Proof of Lemma 2

Recall that expression (2) shows that $\text{Var}(\varepsilon_P|\cdot)$ depends negatively on $\text{E}(\theta|\cdot)$. This proofs, therefore has to show that increases in Z_t increase $\text{E}(\theta|\cdot)$. An increase in Z_t for, $t = 1, 2$, increases the weight of the Beta distribution with the lower mean and thus decreases the mean of the mixture. That is, for any α and β we have that $\text{E}[\text{Be}(\alpha+1, \beta)] = (\alpha+1)/(\alpha+\beta+1) > \text{E}[\text{Be}(\alpha, \beta+1)] = \alpha/(\alpha+\beta+1)$ for the case of Z_1. Given that $p(Z_1)$ is strictly decreasing in Z_1, then it is clear that $\partial\text{E}[\theta|Z_1]/\partial Z_1 < 0$. For the case of Z_2, notice that for any α and β we have that $\text{E}[\text{Be}(\alpha+2, \beta)] > \text{E}[\text{Be}(\alpha+1, \beta+1)] > \text{E}[\text{Be}(\alpha, \beta+2)]$. To complete the argument, notice that for a given Z_1 $\text{E}[\theta|Z_2, Z_1]$ is strictly increasing in the ratio $q(Z_2, Z_1)/r(Z_2, Z_1)$ which is itself decreasing in Z_2. $\square$

Proofs of Propositions 1 to 5 still to be typed in.

References

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

[2] Baker, G., Gibbons, R. and Murphy, K.J. (1994). "Subjective Performance Measures in Optimal Incentive Contracts." *Quarterly Journal of Economics*, 109, pp. 1125-1156.

[3] Bellemare, C. and Sebald, A. (2018). "Self-Confidence and Reactions to Subjective Performance Evaluations." CESifo Working Paper 7325.

[4] Bergemann, D., Bonatti, A. and Smolin A. (2018). "The Design and Price of Information." *American Economic Review*, 108, pp. 1-48.

[5] Esó, P. and Szentes, B. (2007). "The Price of Advice." *RAND Journal of Economics*, 38, pp. 52-54.

[6] Fincham, R. (1999). "The Consultant-Client Relationship: Critical Perspectives on the Management of Organizational Change." *Journal of Management Studies*, 36, pp. 335-351.

[7] Fischer, P. E. and Verrecchia, R. E. (2000). "Reporting Bias." *The Accounting Review*, 75, p. 229.

[8] Fudenberg, D. and Tirole, J. (1986). "A "Signal-Jamming" Theory of Predation." *RAND Journal of Economics*, 17 pp. 366-376.

[9] Gibbs, M., Merchant, K., Van der Stede, W. and Vargus M. (2004). "Determinants and Effects of Subjectivity in Incentives." *The Accounting Review*, 79, pp. 409-436

[10] Holmström, B. (1999). "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies*, 66, pp. 169-182.

[11] Iossa, E. and Jullien, B. (2012). "The Market for Lawyers and Quality Layers in Legal Services" *RAND Journal of Economics*, 43, pp. 677-704.

[12] Ishida, J. (2012). "Contracting with Self-Esteem Concerns." *Journal of Economic Behavior and Organization*, 81, pp. 329-340.

[13] Kaya, A. (2010). "When Does It Pay to Get Informed?," *International Economic Review*, 51, pp. 533-551.

[14] Klein, N. and Mylovanov, T. (2017). "Will Truth out?—An Advisor's Quest to Appear Competent." *Journal of Mathematical Economics*, 72, pp. 112-121.

[15] Lee, P. M. (1997). *Bayesian Statistics*. Arnold (London), Oxford University Press (New York).

[16] Levy, G. (2005). "Careerist Judges and the Appeals Process", *RAND Journal of Economics*, 36, pp. 275-297.

[17] MacLeod W. B. (2003). "Optimal Contracting with Subjective Evaluation." *American Economic Review*, 93, pp. 216-240.

[18] Mailath, G.J. (1987), "Incentive Compatibility in Signaling Games with a Continuum of Types", *Econometrica*, 55, pp. 1349-1365.

[19] Mailath, G.J. and von Thadden, E.-L. (2013), "Incentive Compatibility and Differentiability: New Results and Classic Applications", *Journal of Economic Theory*, 148, pp. 1841-1861.

[20] Mohe, M. and Seidl, D. (2011). "Theorizing the Client-Consultant Relationship from the Perspective of Social-Systems Theory." *Organization*, 18, pp. 3-22.

[21] Prendergast, C. (1993). "A Theory of "Yes Men"." *American Economic Review*, 83, 757-770.

[22] Prendergast, C. and Stole, L. (1996). "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning." *Journal of Political Economy*, 104, pp. 1105-1134.

[23] Sturdy, A. (1997) "The Consultancy Process - An Insecure Business?," *Journal of Management Studies*, 34, pp. 389-413.