Too good to be truthful: Why competent advisers are fired*

Christoph Schottmüller†
University of Copenhagen and Tilec

August 10, 2017

Abstract

A decision maker repeatedly asks an adviser for advice. The adviser is either competent or incompetent and his preferences are not perfectly aligned with the decision maker's preferences. Over time, the decision maker learns about the adviser's type and fires him if he is likely to be incompetent. If the adviser's reputation for being competent improves, it will be more attractive for him to push his own agenda because he is less likely to be fired for incompetence. Consequently, very competent advisers are also fired with positive probability because they pursue their own goals. The quality of advice is highest if the adviser's competence is uncertain.

JEL codes: C73, D83, G24

Keywords: advice, cheap talk, reputation

1. Introduction

As specialization is one of the cornerstones of the modern knowledge society, it is unsurprising that advice given by specialized experts is important in so many domains of life. Savers have financial advisers to help them manage their wealth, consumers rely on sales personnel, politicians and managers depend on their advisers to find the right policy, patients need their physicians' advice and internet users rely on search engines.

In most of these cases the adviser's incentives are not necessarily aligned with the advice seeker's preferences. Financial advisers (as well as sales personnel and search engine operators) can obtain bonuses if their customers buy specific products, while

*I want to thank Ole Jann, Jesper Rüdiger, Peter Norman Sørensen, Juuso Välimäki, audiences at the University of Copenhagen, SING12 in Odense and the Meeting of the EEA in Geneva, two anonymous referees and the associate editor for helpful comments.

†email: christoph.schottmueller@econ.ku.dk, phone: +45 35 32 30 87

politicians and managers might wonder whether their advisers have an own agenda and patients might be worried that their physician's enthusiasm for a certain drug stems from successful lobbying efforts on the part of its producer. Even ex post, it is hard to detect whether these worries were justified because the advice in all these areas is complex and even the best possible advice could turn out to be wrong once in a while.

Another common feature of these examples is the repeated nature of the advice. Most people tend to receive advice from the same adviser several times and switch advisers only from time to time. It clearly makes sense to switch if one concludes that the adviser is not competent i.e. the adviser is more of a quack than an expert. However, long term advisers otherwise viewed as very competent are also occasionally fired. For instance, in 2003, financial analyst Jack Grubman was banned by the Security and Exchange Commission from the financial industry for life and fined fifteen million dollars for misconduct. Grubman had used his good reputation to pursue his personal goals when he gave a public buy recommendation for AT&T as part of a complicated plan for his children to be admitted to the prestigious 92nd Street YM-YWHA's preschool program (as he explained in a private email that later went public).[1] By the time the ban was announced, market participants had, of course, already stopped listening to Grubman's advice. This reaction was, however, not a response to perceived incompetence. When Grubman was hired by Distinctive Devices as consultant a year later, the company's stock price increased. The problem was that Grubman apparently (ab-)used his good reputation by misrepresenting his information and thereby manipulated his followers for his own personal benefit.

To take an example of political advice giving, consider the firing of Roger Stone as Donald Trump's campaign adviser in his race to become the Republican Party's candidate in the 2016 presidential election. Trump and Stone had worked together for more than a decade and Stone was well regarded within the Republican party. Trump explained his dismissal by saying: "I terminated Roger Stone last night because he no longer serves a useful function for my campaign." Trump added: "I really don't want publicity seekers who want to be on magazines or who are out for themselves. This campaign is not about them." In other words, the firing was not due to incompetence but the fact that, from Trump's point of view, Stone prioritized his personal agenda over the one his boss had in mind.

History is full of further examples in which kings have dismissed or even killed their most prominent advisers when these advisers were too competent and perceived as a threat to the throne. Famous in this sense is the Ottoman Sultan Suleiman the Magnificent, who killed not only his Grand Vizier and childhood friend Pargali Ibrahim

[1]See `http://observer.com/2010/03/stockgoosing-grubman-to-sell-townhouse-for-196-m/` for a brief summary of the story. (accessed on August 1, 2017)

Pasha (after Ibrahim committed the mistake of using the title "Serasker Sultan") but also his own son and designated heir Mustafa for this reason (after a successful military campaign on his father's behalf where Mustafa committed the mistake of not stopping his soldiers referring to him as "sultan").

In all these examples, a very competent adviser was mistrusted and fired after committing some "mistake" that made the decision maker doubt whether the adviser acted in the decision maker's best interest or whether he was instead (ab-)using his power to push his own agenda. This paper argues that these situations are typical. More specifically, advisers are fired not *although* they had a reputation for being competent but *because* they had a reputation for competence. That is, they might have kept their positions – and possibly even avoided the mistake – if their competence had been in doubt.

What is the logic behind this result? I consider a setting in which the competence of the adviser is not perfectly known by the decision maker. An adviser whose competence is in doubt is facing the danger of being dismissed for incompetence if his advice turns out to be bad (which will strengthen the decision maker's initial doubts). Consequently, the adviser has strong incentives to act in the decision maker's best interest to keep his position. An adviser who is believed – with high probability – to be competent, however, has more freedom because the risk of him being fired *due to incompetence* in the near future is negligible. That is, even if his advice turns out to be bad a few times, this is not immediately a sign of incompetence as it could simply be due to bad luck. The adviser is therefore free to pursue his own goals, which are usually not in line with the decision maker's goals. Hence, in this case, the best response of the decision maker is to fire the adviser because his advice serves only the interests of the adviser himself and not the decision maker's interests.[2]

Figure 1 shows the reasons for firing an adviser who is, in fact, competent for different beliefs of competence. The decision maker fires the adviser if the belief that he is competent is too low because the information that the adviser is competent is hidden. If the belief is high, then the reason for firing the adviser is not hidden information but moral hazard: the adviser does not act in the interest of the decision maker but pushes his own agenda. Note that the decision maker gets the best advice when he is uncertain about the quality of the adviser because this uncertainty will incentivize the adviser to give good advice.

The model is a repeated game in which the adviser recommends one of two options to

[2]To be more precise, the decision maker and adviser might in equilibrium use mixed strategies when the adviser is believed to be competent. Thus, the decision maker will fire the adviser with some probability when a recommendation turns out to be bad. This gives the adviser some incentive to give not too bad advice to avoid being fired. In equilibrium, the quality of advice will be just high enough to make the decision maker indifferent between firing and keeping the adviser. Otherwise, the decision maker's threat of firing would not be credible.

Figure 1: Reasons for firing a competent adviser (α is the belief that the adviser is competent). The adviser is not fired for intermediate values of α.

the decision maker in every period until the decision maker ends the advice relationship. One of the available options fits the decision maker's needs and one fits the adviser's needs – e.g. he receives a bonus for this option. The two might accidentally coincide from time to time but often they do not. The decision maker has a uniform prior concerning which option will fit his needs and also concerning which will fit the adviser's needs. The adviser receives a noisy signal of which option fits the decision maker's needs and knows perfectly which option will give him a bonus. The decision maker finds out whether the recommended option has fitted his needs only after he has followed the recommendation. The adviser has one of two types: either he is competent – i.e. his noisy signal is informative – or not.

In this model, no meaningful advice could be obtained in a static setting because the adviser would always recommend his bonus option if he did not face the threat of losing future bonus payments. The same is true in a finitely repeated game: similar to the static setting, the adviser is unable to give meaningful advice in the last period and as a consequence, he will always be fired before the last period. Given this, the adviser is unable to give meaningful advice in the second to last period and the game unravels, meaning that the adviser is never consulted in equilibrium. Some informative advice is, however, possible in an infinitely repeated game setting. Unsurprisingly, the adviser is fired for sure if the decision maker's belief that the adviser is competent is very low. If this belief is sufficiently high, then the adviser is also fired with positive probability whenever he recommends an option that does not fit the decision maker's needs. For these high beliefs, equilibrium strategies are usually mixed. The decision maker is indifferent between firing and keeping the adviser and the threat of firing is just high enough to ensure that the adviser finds a strategy optimal that keeps the decision maker indifferent between these two options.

The expected length of the game, i.e. the number of periods before the adviser is fired, is uniformly bounded from above for any belief; that is, the bound is independent of the decision maker's belief about the adviser's competence. This illustrates that even an arbitrarily competent adviser will almost surely be fired within a finite amount of time. These results hold for all equilibria of the game, i.e. they are not affected by multiplicity of equilibria. It is also shown that the adviser suffers in many equilibria from a severe commitment problem: if he was able to commit to truthfully revealing his signal in every period, then he and the decision maker would both obtain a strictly higher payoff than in equilibrium.

Related to this paper is the literature on cheap talk started by Crawford and Sobel (1982) and surveyed in Krishna and Morgan (2008).[3] The exact structure of the payoffs is, however, somewhat different from the traditional cheap talk setup where the adviser has a bias in a certain direction, e.g. a political adviser is more left-wing than the decision maker and will therefore always push for more leftist policies than the decision maker would like. In my setup, the bias does not run in a certain direction but is for one (random) option which carries a bonus for the adviser. Within the cheap talk literature, models of repeated cheap talk in which the state of the world changes each period are closest to my paper.[4] Renault et al. (2013) characterize the set of equilibrium payoffs in a repeated game framework when players are arbitrarily patient and states are correlated through an irreducible Markov chain. Park (2005) analyzes a situation where a consumer has a problem each period and relies upon advice to find out which of the several repair shops specializes in fixing the problem at hand. In contrast to the current paper, the adviser in these papers knows the state of the world perfectly; consequently, reputation for competence does not play a role. Hence, an expert cannot be "too good" which is the main focus of my paper. This is also the main difference to earlier papers (Sobel, 1985; Benabou and Laroque, 1992) where the adviser's type refers to his honesty and not the quality of his information. Being more certain of facing a honest type is good for consumers, implying that they will certainly not fire the adviser at these favorable beliefs. In my model, being more likely to face an informed type can be bad because it aggravates the moral hazard problem.

An effect closely related to this paper is known from the literature on reputation with repeated interaction between a sequence of short lived principals and a long lived agent where both moral hazard and adverse selection are present, see Mailath and Samuelson (2001), Jullien and Park (2014). In these papers, incentivizing agents with a good reputation to exert high effort becomes so difficult that equilibria with high effort at some reputation level do not exist (if monitoring is noisy).[5] The reason is – similar to my paper – that the belief updating is slow when the belief of facing a competent type

[3]The adviser's advice is in my setup directly relevant for the decision maker's payoff but the main reason is that – due to his ignorance – the decision maker has no real choice but to follow the adviser's advice (as long as he did not quit the advice relationship). That is, the model of this paper is equivalent to a model where the advice is real cheap talk and the decision maker has a pseudo decision to follow the advice or not.

[4]There is also a literature analyzing the effect of repeated advice when the state is the same in all periods and only one action has to be taken, e.g. Aumann and Hart (2003), or an action is taken repeatedly, e.g. Golosov et al. (2014).

[5]One exception is proposition 2 in Mailath and Samuelson (2001) where a mixed equilibrium is constructed in which high effort is exerted for some reputation levels. However, the set of these reputation levels is countable and strategies necessarily have an infinite number of discontinuities in any such equilibrium; see section 5 for a more detailed discussion.

is close to 1. However, a high effort equilibrium may exist if types are impermanent; that is, in each period the type of the agent changes – unobserved by the principal – with a small probability. A similar logic also appears in Cripps et al. (2004) and Wiseman (2008). My paper differs in several ways: when looking at Markov equilibria, the decision maker's strategy is allowed to depend not only on "reputation" – that is the belief that he is competent – but also on whether advice in the current period was successful (while it depends only on reputation in the aforementioned papers). This allows for partially informative equilibria that still display the difficulty of incentivizing advisers with high reputation: advisers with a good reputation give worse advice and are fired with some probability. In contrast, Mailath and Samuelson (2001) and Jullien and Park (2014) focus on "honest Markov equilibria" in which good agents exert high effort regardless of their reputation and a higher reputation is therefore always good news for the principal. Furthermore, types are permanent in my model and both adviser and decision maker are long lived and strategic (which would lead to a complete breakdown of communication in the papers mentioned above). Having a long lived decision maker allows a focus on stopping/firing which is less relevant in the aforementioned papers with short lived principals. In addition, several of my results apply to all perfect Bayesian equilibria and not only Markov equilibria.

More broadly, the paper is part of the literature asking whether career concerns and reputation can prevent opportunism, see Fama (1980) and Holmström (1982) for seminal contributions. Closest is Aghion and Jackson (2016) in which (political) "leaders" have to be incentivized to take risky decisions (instead of remaining inactive) by the threat to vote them out of office. In equilibrium, even arbitrarily competent leaders are terminated with some probability whenever they do not take a risky decision. However, setup and applications differ significantly as "leaders" do not receive bonuses and do not know their own type.

Another strand of the literature, e.g. Brandenburger and Polak (1996); Ottaviani and Sørensen (2006a,b), analyzes how an expert who wants to maximize his reputation for being competent will misrepresent his information. The main result is that the adviser will then misrepresent his signal towards the prior. This interplay between prior beliefs and reputation concerns in a repeated game setup is also present in many other papers, e.g. Prendergast and Stole (1996), Morris (2001), Ely and Välimäki (2003), Li (2007) and Klein and Mylovanov (2016). This effect is not present in the current paper as the decision maker will have a uniform prior which makes it impossible to misrepresent towards the prior. Furthermore, the expert wants to maximize his expected bonus stream – and not his reputation per se – which leads to the aforementioned moral hazard problem that drives the results of my paper.

The outline of the paper is as follows. The next section introduces the model and describes the solution concepts used. Section 3 presents the results – most prominently that competent advisers are fired with positive probability and that the game is expected to end within a given finite time. It also gives necessary and sufficient conditions for the existence of Markov equilibria with informative communication and points out a commitment problem on the side of the adviser. The model is extended to allow for monetary transfers and some competition among several advisers in section 4. Section 5 discusses the results and their implications and section 6 concludes. Proofs are relegated to the appendix.

2. Model

Actions and payoffs: As long as the decision maker (DM) does not fire the adviser (A), the stage game in period t is as follows: A receives a noisy signal about DM's needs in this period and recommends one of two available options to DM. Only one of the two options fits DM's needs and DM receives a payoff of 1 (0) if the recommended option fits (does not fit) his needs.[6] DM's payoff is observed by both players, i.e. both players observe whether the recommendation fitted DM's needs or not. One of the two options leads to a "bonus" for A. That is, A receives a payoff of 1 (0) if he recommends (does not recommend) the bonus option. The identity of the bonus option and A's payoff are privately observed by A. At the end of the period, DM decides whether to continue to the next period or to stop the game, i.e. fire the adviser, which gives him an outside option payoff of W_O in $t + 1$ and zero payoff thereafter. A's payoff is zero in all future periods if DM stops the game. Both players discount future payoffs using the common discount factor $\delta \in (0, 1)$. The timing of the stage game is summarized in figure 2.

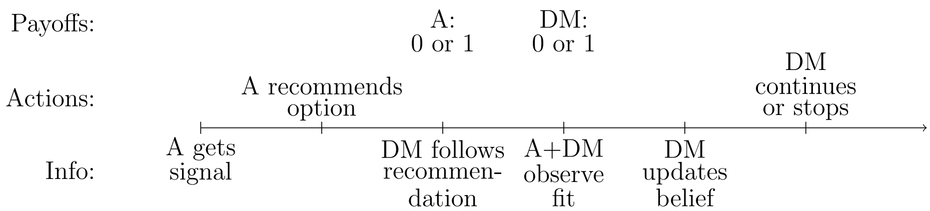

Figure 2: Timeline within a given period t.

Information: Each option fits DM's needs with probability 1/2 and is A's bonus option with probability 1/2. These events are independent. This implies that DM cannot infer from the identity of the recommended option the likelihood that A recommends the bonus option.[7] DM's needs and A's bonus options are uncorrelated across periods.

[6]The recommendation is directly payoff relevant. However, the model is equivalent to one where DM has a pseudo decision of choosing an option and this decision is payoff relevant.

[7]One interpretation is that DM only knows that there are 2 options but is unaware of what these options are; that is, he learns that a specific option exists only if it is recommended to him (this is

(These assumptions are made for ease of exposition and can be relaxed, see section 3.1.)

A has one of two types. If A is *incompetent*, his signal about DM's needs is completely uninformative. An incompetent type's belief assigns therefore probability 1/2 to each option fitting DM's needs. If A is *competent*, his signal is noisy but informative. More precisely, a competent type's belief concerning DM's needs assigns probability $p \in (1/2, 1)$ to one option and $1 - p$ to the other option. I will refer to the more likely option as "option h" ($h$igh probability fit) and the less likely option as "option l" in the remainder (but keep in mind that DM is not aware of these identities as A's signal is private!). For simplicity, p is assumed to be time invariant. A's type is A's private information while also being time invariant.

DM's belief as to whether A is competent is denoted by $\alpha \in [0, 1]$. This belief is updated using Bayes' rule after observing whether the recommendation has fitted DM's needs or not. DM's initial belief α_0 is assumed to be common knowledge. Consequently, α will be commonly known in equilibrium as both players observe whether the recommendation has fitted DM's needs. I will occasionally refer to α as *A's reputation.*

Parameter assumption: To make the problem interesting, DM's outside option is assumed to satisfy the inequalities

$$\frac{1/2}{1-\delta} < W_O < \frac{p}{1-\delta}. \tag{1}$$

That is, DM prefers his outside option to choosing an option randomly for each period, but he prefers choosing an option fitting his needs with probability p each period to his outside option. Consequently, DM would like to fire an incompetent adviser, but would prefer to keep a competent adviser if this adviser reveals his signal truthfully in every period.

2.1. Strategies, Value Functions and Equilibrium

I use perfect Bayesian equilibrium (referred to as "equilibrium" from here onward) as main solution concept. In any such equilibrium the incompetent type will always recommend his bonus option: the incompetent type cannot influence the probability with which DM's needs are satisfied because both options are equally likely to do so from the point of view of the incompetent type. Consequently, recommending his bonus option has – in expectation – the same consequences for future play as recommending the other option, but the bonus option gives an immediate payoff of 1.

the reason why he seeks advice in the first place). Think of a person googling "Italian restaurants in Manhattan" or a patient asking a physician for the right medication. Another example would be that customers in financial advice usually do not know the bonuses that are associated with all possible investments. In these cases, advice is an experience good.

Lemma 1. *In every perfect Bayesian equilibrium, the incompetent adviser recommends the bonus option in every period (regardless of history).*

As the incompetent type's behavior does not require any further analysis I will refer with "adviser" (A) to the competent type in the remainder of the analysis.

For some of the results, I will focus on Markov equilibrium as solution concept which is a common refinement of perfect Bayesian equilibrium. Markov strategies depend only on the state variable which is the adviser's reputation α and observations in the current period. In other words, a *Markov strategy* for A depends on the identity of the bonus option and the belief α. A Markov strategy for DM depends on whether the recommendation fitted DM's needs in this period and the belief α. A mixed strategy of DM is denoted by two measurable functions (β^+, β^-) where $\beta^+ : [0,1] \to [0,1]$ assigns to each α the probability with which DM continues the game if the recommendation of the current period has fitted his needs. $\beta^- : [0,1] \to [0,1]$ assigns to each belief α the probability with which DM continues if the recommendation did not fit DM's needs. Note that these two functions will not be identical in general. Put differently, the probability of stopping the game in a given period t will not only depend on the reputation but also on whether the recommendation in t fits DM's needs or not. It is convenient to let β^+ and β^- depend on the updated belief (after observing whether the recommendation has fitted DM's needs) and I will follow this convention.[8] A mixed Markov strategy of A can be written as $s : [0,1] \times \{l,h\} \to \Delta\{l,h\}$ where $s(\alpha, b)$ denotes the probability distribution over recommended options if A's reputation is α and his bonus option is $b \in \{l,h\}$. A profile of Markov strategies that constitutes a perfect Bayesian equilibrium is called "Markov equilibrium".

A's value function in a given Markov equilibrium is denoted by $V : [0,1] \to \mathbb{R}_+$. That is, $V(\alpha)$ denotes A's expected discounted payoff stream at the very start of a period, i.e. before observing the identity of the bonus option, if A has reputation α. Similarly, DM's value function is denoted by $W : [0,1] \to \mathbb{R}_+$. For some initial belief α I denote the updated belief in case the recommendation has (not) fitted DM's needs by α^+ (α^-). When using Markov equilibrium as a solution concept I will restrict myself to "informative" Markov equilibria which I define as Markov equilibria in which $\alpha^+ \geq \alpha^-$ for all $\alpha \in [0,1]$.[9] That is, the competent type is at least as likely to recommend the option fitting DM's needs as the incompetent type and consequently a recommendation

[8]Note that this is equivalent to having these functions depend on the belief at the beginning of a period because DM's strategy depends directly on whether the recommendation has fitted his needs or not.

[9]The restriction to *informative* Markov equilibria is without loss of generality if – instead of having directly payoff relevant recommendations – DM is allowed to choose an option himself after receiving the recommendation. The reason is that DM would simply choose the option that was not recommended if the competent type gave "worse" advice than the incompetent type.

fitting DM's needs will not worsen A's reputation.[10] From here onward, I will refer to informative Markov equilibrium simply as *Markov equilibrium*.

The restriction to informative Markov equilibria leads to the following technical result which states that A's expected future payoff stream is more valuable after a fitting than a non-fitting recommendation.

Lemma 2. *Let V be the value function in a Markov equilibrium. Then, $\beta^+(\alpha^+)V(\alpha^+) \geq \beta^-(\alpha^-)V(\alpha^-)$.*

A's expected utility in a given period can be written as

$$q\,\delta\beta^+(\alpha^+)V(\alpha^+) + \delta(1-q)\beta^-(\alpha^-)V(\alpha^-) + \mathbb{1}_{bonus}$$

where q is the probability that the recommendation satisfies DM's needs (which depends on the specific recommendation) and $\mathbb{1}_{bonus}$ is the indicator function for the bonus option, i.e. it is 1 if A recommends the bonus option and 0 otherwise. As lemma 2 states that $\beta^+(\alpha^+)V(\alpha^+) \geq \beta^-(\alpha^-)V(\alpha^-)$, it is clear that A will always recommend option h if option h is the bonus option: This recommendation maximizes the chance of improving his reputation and also pays him a bonus. If option l is the bonus option, however, then A might recommend either option depending on the exact values of p, $\beta^+(\alpha^+)V(\alpha^+)$ and $\beta^-(\alpha^-)V(\alpha^-)$. That is, A will only recommend option h if expected future benefits in form of reputation gains from good advice offset the foregone bonus. The following lemma states this more formally.

Lemma 3. *Let V be A's value function in a Markov equilibrium. A's strategy in Markov equilibrium satisfies the following:*

1. *if option h is the bonus option, A recommends option h;*
2. *if option l is the bonus option, A recommends option h only if*

$$\beta^+(\alpha^+)V(\alpha^+) - \beta^-(\alpha^-)V(\alpha^-) \geq \frac{1}{\delta(2p-1)}; \tag{2}$$

and recommends option l only if the reverse inequality holds.

A is willing to mix if and only if (2) holds with equality.

DM's optimal strategy is relatively simple: He ends the game if his expected payoff (in $t+1$ and following periods) from continuing the game is lower than his outside option

[10] While non-informative equilibria are economically somewhat nonsensical, they could in principle exist because a competent A has the ability to give worse advice than an incompetent A. If DM expects A to do so (at some point of the game), then it might be a best response for A to give bad advice (at this point of the game) to improve (!) his reputation: fitting advice would then be interpreted as being more likely to be given by an incompetent type and therefore reduce α.

W_O. If the competent type recommends the option fitting DM's needs with probability q (given reputation α), then DM's value is

$$W(\alpha) = (\alpha q + (1-\alpha)/2)(1 + \delta W(\alpha^+)) + (\alpha(1-q) + (1-\alpha)/2)\delta W(\alpha^-). \quad (3)$$

As long as $W(\alpha) \geq W_O$, it is optimal to continue. DM is willing to use a mixed strategy if and only if $W(\alpha) = W_O$ and stops the game if $W(\alpha) < W_O$.

For future reference, denote – for a given initial belief α – the updated beliefs if A always recommends option h (regardless of the identity of the bonus option) by α_h^+ and α_h^-:

$$\alpha_h^+ = \frac{\alpha p}{(1-\alpha)/2 + \alpha p} \quad (4)$$

$$\alpha_h^- = \frac{\alpha(1-p)}{(1-\alpha)/2 + \alpha(1-p)}. \quad (5)$$

If A always recommends the bonus option (even if it is option l), then A behaves exactly like the incompetent type. This means that DM's posterior belief will equal his prior belief, i.e. $\alpha^+ = \alpha = \alpha^-$. The updated beliefs will be between α and the ones in (4) and (5) if A mixes between always recommending option h and always recommending the bonus option.

3. Results

If the adviser's reputation is sufficiently low, it is optimal for DM to stop the game. This follows almost directly from the assumption that DM's outside option is strictly better than receiving advice from an incompetent adviser forever, see (1). If DM is almost sure to face an incompetent adviser, it will therefore be optimal for him to fire the adviser.

Proposition 1. *In equilibrium, there exists an $\underline{\alpha} > 0$ such that DM ends the game whenever $\alpha < \underline{\alpha}$.*

The more interesting result is that DM will also end the game (with some probability) if the adviser's reputation is sufficiently high. The intuition for this result is as follows: Suppose DM continued for sure if α is above some threshold $\tilde{\alpha} < 1$. For α close enough to 1, A would then be very sure that DM would continue even if he gave repeatedly bad advice. This is true as α^- is very close to α if α is close to 1, see (5). Put differently, A has hardly any dynamic incentives to give good advice. Statically, however, he has an incentive to recommend the bonus option as this gives an immediate payoff of 1. A will therefore recommend the bonus option no matter what his signal is. Consequently, both adviser types behave in the same way which has two implications:

First, the belief updating stops, i.e. $\alpha = \alpha^+ = \alpha^-$. Second, DM's expected payoff is below his outside option as this situation gives him the same payoff as receiving advice from an incompetent adviser forever. Clearly, this contradicts our starting point that DM continues whenever $\alpha > \tilde{\alpha}$. As such an $\tilde{\alpha} < 1$ does not exist, we can conclude that there are beliefs arbitrarily close to 1 where DM quits the game with positive probability. Recall that DM ends the game only if his continuation value is (weakly) less than his outside option. Consequently, DM's continuation value will equal W_O for some beliefs α arbitrarily close to 1. Having an adviser with high reputation is therefore not necessarily valuable. The following theorem states this result more formally and strengthens it for Markov equilibria in which both players choose piecewise continuous strategies.[11]

Theorem 1. *In every Markov equilibrium there exists an $\varepsilon > 0$ such that there is a sequence of beliefs $(\alpha_i)_{i=1}^{\infty}$ converging to 1 where DM ends the game with at least ε probability for every element of the sequence.*

If Markov equilibrium strategies are piecewise continuous, then there exists an $\bar{\alpha} < 1$ such that $W(\alpha) = W_O$ for all $\alpha > \bar{\alpha}$. Furthermore, there exists an $\tilde{\varepsilon} > 0$ such that DM continues with probability $\beta^-(\alpha^-) < 1 - \tilde{\varepsilon}$ in case the recommendation does not fit his needs for all $\alpha > \bar{\alpha}$.

The previous theorem established that DM will fire the adviser with positive probability for high α. However, if this probability is close to zero, one could argue that it has little economic relevance. The intuition given above should already illustrate that this is not the case because similar problems as for zero quitting probability also emerge with very small positive quitting probabilities. The following lemma strengthens this intuition by stating that DM quits the advice relationship almost certainly within T periods – where T is some finite number depending on the parameters – no matter what the current belief is. Note that in the following lemma the upper bound on the length of the advice relationship neither depends on the (initial) belief α nor on the equilibrium.

Lemma 4. *Let $\varepsilon > 0$ and define*[12]

$$T_\varepsilon = \left\lceil \frac{\log(\varepsilon)}{\log\left(1 - (1-p)^{T'}\varepsilon'\right)} \right\rceil T' \qquad \textit{where} \quad \varepsilon' = \frac{1-\delta}{2} \quad \textit{and} \quad T' = \left\lceil 2\frac{\log(1-\delta)}{\log(\delta)} - 1 \right\rceil.$$

The probability that DM ends the game within T_ε periods is at least $1 - \varepsilon$ in every equilibrium.

[11]It is natural to state this result for Markov equilibria because it is directly based on the belief α. However, the logic above clearly implies a related result for all perfect Bayesian equilibria: there is no equilibrium in which DM continues with probability 1 whenever his belief is above some threshold $\tilde{\alpha} < 1$; see proposition 3 for a formal result along these lines.

[12]The ceiling $\lceil x \rceil$ is the smallest integer above x.

The previous lemma is based on the following idea. If it is sufficiently unlikely that A is fired – in case of (repeated) bad advice – in the following T periods, then A finds it optimal to recommend his bonus option even if it is option l. However, in this case DM should better fire A right away. Lemma 4 has a direct implication on the expected length of the relationship. Again, the result holds for every (initial) belief α and every equilibrium.

Theorem 2. *The expected length of the advice relationship in equilibrium is finite and bounded from above by*

$$\bar{T} = T' \left(2 - \frac{1}{\log(1 - (1-p)^{T'}(1-\delta)/2)} \right).$$

Theorem 2 is driven by the fact that DM fires the adviser when his reputation is high and not just by the possibility that the adviser is fired due to incompetence as in proposition 1. To see this, consider $\alpha \to 1$. For beliefs arbitrarily close to 1, the time until which the belief α could possibly fall below the incompetence threshold $\underline{\alpha}$ is going towards infinity. Nevertheless, the expected length of the game is below $\bar{T}$ for any belief. That is, the finiteness of the expected game length is driven by DM ending the game for high beliefs. For low beliefs, the result is, of course, driven by the fact that DM fires the adviser if his reputation is too low.

The previous results derived properties of (Markov) equilibria without ensuring the existence of such equilibria. Similar to normal cheap talk models, an equilibrium in which no meaningful advice is given ("babbling equilibrium") will always exist. In the framework of this paper, the babbling equilibrium takes the following form: DM always ends the game and A always recommends his bonus option. As usual in cheap talk, equilibria with (some) information transmission may also exist. The following result states a necessary and sufficient condition for the existence of Markov equilibria in piecewise continuous strategies in which some information is transmitted. I will refer to such equilibria as *communication equilibria*.

Definition 1. *A* communication equilibrium *is a Markov equilibrium in piecewise continuous strategies such that (i) for some interval of beliefs (α_1, α_2) A recommends option h with strictly positive probability even if it is not the bonus option, i.e. $s(\alpha, l) > 0$ for $\alpha \in (\alpha_1, \alpha_2)$, and (ii) for some interval of beliefs (α_3, α_4) DM continues with strictly positive probability, i.e. $\beta^+(\alpha) + \beta^-(\alpha) > 0$ for $\alpha \in (\alpha_3, \alpha_4)$.*

Proposition 2. *If*

$$p \geq \frac{3}{4} + \frac{1-\delta}{2\delta}, \tag{6}$$

then there exists a communication equilibrium in which $s(\alpha, l) > 0$ and $\beta^+(\alpha) > 0$ for *all beliefs*

$$\alpha > \underline{\alpha} = \frac{W_O(1-\delta) - 1/2}{p - 1/2}.$$

If there exists a communication equilibrium, then (6) holds and

$$\lim_{\alpha \to 1} V(\alpha) = \left[\frac{4p-1}{4p-2}, \frac{4p-3}{(4p-2)(1-\delta)}\right]. \tag{7}$$

Condition (6) holds if A's signal is sufficiently informative and players are sufficiently patient. In particular, $p > 3/4$ and $\delta > 2/3$ are both implied by (6). This condition intuitively makes sense in so far as to incentivize A to give useful recommendations instead of cashing in on his bonus immediately, A must enjoy substantial payoffs in the future. A cannot be incentivized by future payoffs if he discounts those a lot. If the signal technology is bad, the welfare gains from advice are small. DM has to leave some of these welfare gains to the adviser to incentivize him; some rents go to the incompetent type and sometimes A will recommend his bonus option even if it is option l (theorem 1 implies that this has to be true in every communication equilibrium). If the signal technology is too bad, the payoff left for DM is simply too small to prevent him from taking his outside option.

The last part of proposition 2 illustrates a basic commitment problem A faces. Suppose A could commit to the strategy "always recommend option h". This would give DM the highest possible payoff and imply that DM does not stop the game if he believes that A is sufficiently competent. What is more surprising is that this commitment would also increase the payoff of A. Note that the probability of recommending the bonus option would be 1/2 in each period entailing that the expected payoff of A under commitment would be $1/(2(1-\delta))$. It is straightforward to verify that this payoff is higher than the upper bound of $\lim_{\alpha \to 1} V(\alpha)$ in (7).

Corollary 1. *A's payoff in a communication equilibrium is – for high values of reputation – lower than the commitment payoff of a competent type:* $\lim_{\alpha \to 1} V(\alpha) < 1/(2(1-\delta))$.

3.1. Relaxing Assumptions

In this section, I want to discuss how some of the assumptions could be relaxed. First, the assumption that A and DM share the same discount factor can be discarded without affecting any result or proof as long as the two discount factors are strictly less than 1. Second, one might wonder about a signal technology that is not constant over time; that is, the posterior p of a competent type might depend on the specific period. Strategies should then naturally depend on the time period as well, implying that Markovian strategies make less sense in this setting. The results on the (expected) length of the game, however, still hold true if one substitutes p by $\sup(\{p^t\})$ where p^t is the competent type's posterior in period t and the supremum is assumed to be strictly below 1. With this adjustment the proofs of lemma 4 and 2 will go through and the results will hold. Third, one can allow for more than 2 options at the expense of a somewhat more cluttered notation; see an earlier working paper version of this paper (Schottmüller, 2016).

Finally, I want to discuss the possibility of correlation between A's bonus option and the option fitting DM's needs. First, consider within period correlation of the competent type's bonus option with DM's needs. As long as the correlation is imperfect, this will not change theorems 1 and 2 qualitatively because the only relevant part of the signal technology is A's posterior belief which was denoted by p. Correlation will now imply that the posterior p can depend on the identity of the bonus option. As long as the posterior is strictly between 1/2 and 1, the structure of the problem does not change and the proofs go through using the higher of the possible posteriors p. (Of course, also the updating of DM's belief α will be affected but this does not change the results.) Second, consider correlation over time. More precisely, the probability that A's bonus option coincides with the option fitting DM's needs might depend on whether it had also done so in the previous period.[13] Note that this form of correlation is irrelevant for A's decision problem because he knows his bonus option when making his recommendation (and the correlation does not improve the accuracy of his signal). The correlation will affect DM's beliefs and decisions. The proofs of theorems 1 and 2 (and lemma 4) are based on scenarios where A wants to recommend his bonus option although it is option l. Hence, as long as this scenario has positive probability – that is, so long as the correlation is not perfect – the same situation emerges and the results hold qualitatively.

One might also wonder whether it is possible to restrict the strategy space such that DM's decision will depend upon his belief α only (and not on whether the recommendation in the current period fitted his needs or not). Put differently, could one impose $\beta^+ = \beta^-$? Unfortunately, no communication equilibrium exists in this restricted class. This result is obtained by Mailath and Samuelson (2001) in a very similar model. Their proposition 2.3 implies that strategies in a Markov equilibrium using this restricted strategy space cannot be piecewise continuous (unless A babbles). They construct an equilibrium with an infinite number of discontinuities (proposition 2.2). However, the adviser's strategy in this equilibrium is equivalent to a babbling equilibrium for all but a countable number of beliefs α. I will briefly sketch the argument why no Markov equilibrium with piecewise continuous strategies apart from the babbling equilibrium exists. Consider $\beta(\alpha^+)V(\alpha^+) - \beta(\alpha^-)V(\alpha^-)$ as $\alpha \to 1$. With piecewise continuous strategies V and β will be continuous for all α above a certain threshold. From Bayes'

[13] Here I still maintain the assumption that at the start of the period DM views both options as equally likely to fit his needs. Otherwise, the strategic effects discussed, for example, in Ottaviani and Sørensen (2006a,b) will play a role.

rule, see (4) and (5), it then follows that $\alpha^+ - \alpha^- \to 0$ as $\alpha \to 1$. Together with the continuity of V and β, this implies that $\beta(\alpha^+)V(\alpha^+) - \beta(\alpha^-)V(\alpha^-)$ converge to 0 as $\alpha \to 1$. Consequently, A will recommend his bonus option by lemma 3 for beliefs above some threshold $\hat{\alpha} < 1$. Now suppose that – for some lower belief – A recommended option h with positive probability if option l is his bonus option. Looking at the highest such belief α, it is evident that DM should stop the game after a fitting recommendation as (i) $\alpha^+ > \alpha$ and (ii) A always recommends his bonus option for all beliefs above α. But in this case it is certainly not optimal for A to recommend option h if option l is his bonus option and his reputation is α which contradicts the fact that he does so with positive probability. Hence, no such α exists and babbling is the only Markov equilibrium. The intuitive assumption that DM's equilibrium strategy depends on whether he just got fitting advice (or not) is therefore essential for constructing communication equilibria.

4. Extensions

4.1. Transferable utility

So far, DM could not use monetary payments to incentivize the agent. One might conjecture that using monetary rewards for fitting advice could mitigate the moral hazard problem that leads to the firing of advisers with a high reputation. This section shows that theorems 1 and 2 continue to hold if monetary transfers from DM to A are allowed.

The model now allows DM to transfer a wage w to A at the end of each period t; i.e. at the same time when DM decides whether to continue or to stop. (The exact timing of the payments does not matter for the results.) A's payoff in a given period is $1 + w$ if he recommends his bonus option and w else where w is the wage payment he receives. DM's payoff in a given period is $1-w$ if the recommendation fits his needs and $-w$ else. In the spirit of relational contracting, I will not allow the players to commit to an enforceable payment schedule but will consider self enforcing transfers only.[14] I will restrict the analysis of this section to *informative equilibria*, i.e. perfect Bayesian equilibria in which $\alpha^+ \geq \alpha \geq \alpha^-$ for all $\alpha \in [0, 1]$ and all subgames.

It then remains true that DM must fire A with some probability for some arbitrarily high beliefs. The intuition here is similar to the one presented previously: Suppose DM never fired the adviser if his belief is above some thereshold $\bar{\alpha} < 1$. For every $T \in \mathbb{N}$ one can then find a belief α_T sufficiently close to 1 such that an adviser with reputation α_T will maintain a reputation above $\bar{\alpha}$ even if A's recommendations do not

[14] With commitment, the usual way of solving moral hazard problems – selling the enterprise to the agent – will work. In the context of this paper then, this would entail that DM offers A to commit to paying him 1 unit for every fitting recommendation if A pays $p/(1-\delta) - \varepsilon$ to DM ex ante. The competent type will accept the offer and the incompetent will not.

fit DM's needs for T consecutive periods. An adviser with reputation α_T (or higher) can therefore guarantee himself a payoff of $\sum_{t=0}^{T} \delta^t$ by recommending his bonus option T times. Hence, his equilibrium payoff must be at least $\sum_{t=0}^{T} \delta^t$. For sufficiently high T one can show that $\sum_{t=0}^{T} \delta^t$ is higher than the welfare the advice relationship generates. Consequently, DM would be obliged to have an expected payoff below the payoff he would obtain by immediately stopping the game which is impossible in equilibrium.

To state the following results define $\eta = (1-\delta)W_O - 1/2$ and note that $\eta > 0$ by (1).

Proposition 3. *In an informative equilibrium, there exists no $\bar{\alpha} < 1$ such that DM ends the game with probability less than $\varepsilon = (\eta - \eta\delta)/(2 - 2\eta\delta) > 0$ in all subgames in which his belief is above $\bar{\alpha}$.*

Similar results as in lemma 4 and theorem 2 also still hold.

Proposition 4. *Let $\varepsilon > 0$ and define*

$$T_\varepsilon = \left\lceil \frac{\log(\varepsilon)}{\log\left(1-(1-p)^{T'}\varepsilon'\right)} \right\rceil T' \qquad \textit{where} \quad \varepsilon' = \frac{\eta(1-\delta)}{2-2\delta\eta} \quad \textit{and} \quad T' = \left\lceil \frac{\log(\eta/2)}{\log(\delta)} \right\rceil.$$

The probability that DM ends the game within T_ε periods is at least $1-\varepsilon$ in every informative equilibrium.
The expected length of the advice relationship in equilibrium is finite and bounded from above by

$$\bar{T} = T'\left(2 - \frac{1}{\log(1-(1-p)^{T'}\varepsilon')}\right).$$

There are two caveats which should be made to the results above. First, the results do not hold if A can also make payments to DM. In this case, there are equilibria in which both types of the adviser will compensate DM for bad recommendations and DM will continue forever, regardless of which type he is facing. Such equilibria exist because, from a total welfare point of view, it is efficient for DM to interact even with an incompetent type: the incompetent type's bonus payments are more valuable than DM's outside option by (1). It is then unsurprising that the efficient outcome can be obtained with payments from A to DM if players are sufficiently patient.

Second, to derive the results it was assumed that A's gain from recommending his bonus option was as high as DM's payoff from receiving a fitting recommendation – both were assumed to be 1. Equilibria in which A recommends option h even when his reputation is high can be sustained if A's payoff from recommending the bonus option is sufficiently small. Intuitively, A could be incentivized to always recommend option h if he received an incentive payment of β/p for fitting recommendations where β denotes the payoff that A can get from recommending his bonus option. For β sufficiently small, DM would be willing to make this incentive payment to obtain the best possible advice from a competent type. Consequently, the results above only hold if β is not too small.

4.2. Competition among advisers

One can easily imagine situations in which a decision maker has to decide which of several advisers to consult. In this vein, DM's outside option in the model could be interpreted as the value he assigns to obtain advice from the next best adviser. This section briefly discusses three settings that reflect this idea and shows that the results of the paper carry over to these alternative model settings.

First, assume that there is a finite number of potential advisers – for simplicity, say two. DM's initial belief α_0^i that adviser i is competent might differ between these two advisers. Both advisers have the same preferences and signal technology as in the main model and their types are assumed to be uncorrelated. Furthermore, DM can ask only one adviser each period. First, assume that DM still has an exogenous outside option satisfying (1). At the end of a period DM has then three actions: *stop* the game and take the outside option, *stay* with the current adviser or *switch* to the other adviser.

In this setting, there are direct analogues of theorems 1 and 2. The reason is that both results were proven by showing that A would prefer to recommend his bonus option with probability 1 if the conditions of the theorems were not satisfied (and in this case DM prefers to stop). Note that A is even more tempted to recommend his bonus option in the current setting because there is a possibility that instead of stopping the game, DM will switch and then switch back later. Consequently, the proofs still apply. Theorem 1 implies that for a fixed α^{-i}, there is no $\bar{\alpha}$ such that DM always chooses adviser i if $\alpha^i > \bar{\alpha}$. Since this holds for any given and fixed α^{-i}, the statement is also true without conditioning on α^{-i}: there is no $\bar{\alpha}$ such that DM chooses adviser i whenever $\alpha^i > \bar{\alpha}$. Similarly, lemma 4 and theorem 2 still hold but state now that DM will end the game *or switch* (possibly to an adviser with a worse reputation) within the given number of periods.

As a second setting, consider the same setting as above, but eliminate the exogenous outside option. That is, DM has only the choice between *staying* with the current adviser or *switching* and DM's outside option is therefore completely endogenized. Note that this game still has a babbling equilibrium in which both advisers will always recommend their bonus option and DM's strategy is independent of whether the recommendations have fitted his needs or not, for example, DM always stays with adviser 1. In a literal sense, this babbling equilibrium violates theorem 1 as DM always stays with adviser 1. However, the following analogue to theorem 1 holds trivially. If in a given equilibrium there exists an $\bar{\alpha}$ such that DM stays with adviser i with probability 1 for all $\alpha^i > \bar{\alpha}$, then babbling emerges in all subgames with $\alpha^i > \bar{\alpha}$.[15] A similar analogue

[15]Following the argument in the proof of proposition 3, this result could be strengthened with regard

to theorem 2 would state that the adviser will always recommend his bonus option if the expected length of the advice relationship is above $\bar{T}$ in a given subgame.

Third, I want to consider the possibility that DM has the options to either *continue* or to *reset*. Reset means that DM starts with a new adviser of a given reputation α_0. That is, the outside option W_O is again endogenous. In contrast to the previous setting, DM cannot return to advisers he previously fired. The following result is similar to the sufficiency part of proposition 2.

Proposition 5. *A communication equilibrium exists if (6) holds.*

Furthermore, a similar line of thoughts as earlier shows the same results as the second setting: A will always recommend his bonus option for sufficiently high α if there is an $\bar{\alpha}$ such that DM continues for sure whenever $\alpha > \bar{\alpha}$. A will always recommend his bonus option if the expected length of the advice relationship with the current adviser exceeds $\bar{T}$ as given in theorem 2. Put differently, theorems 1 and 2 hold unless there is babbling.

5. Discussion

The results of this paper establish an inefficiency: the game is expected to end in finite time although the advice relationship lasts forever in a first best world. The inefficiency arises due to the assumption that A cannot commit to a strategy, e.g. the strategy "always recommend the option most likely to fit DM's needs", because his signal is private. Of course, the private nature of A's signal captures exactly the reason why DM needs to get advice. It is perhaps unsurprising that, for example, a consumer may buy the wrong investment products if his financial adviser receives a bonus for selling certain products. The inefficiency established in this paper is, however, of a somewhat subtler nature. Here, not only will the consumer purchase the wrong products, he will also switch to worse financial advisers after some time or not take any advice – depending on how the outside option in the current model is interpreted.

One possible solution to the problem would be to resolve the underlying difference in objectives; that is, to eliminate A's bonus payments. This idea was, for example, expressed in the "Global Analyst Research Settlements" (2003) between US regulators and 10 top investment banks in the aftermath of the dot com bubble. This settlement required banks to separate research and investment banking and stated that the compensation of analysts cannot depend on investment banking activities.[16] The motivation

to the case where DM stays with adviser i with probability of at least $1-(n-n\delta)/(2-2\eta\delta)$ whenever $\alpha^i > \bar{\alpha}$. As in the proof of proposition 3, this would imply that adviser i will always recommend his bonus option for sufficiently high α.

[16]See `https://www.sec.gov/news/speech/factsheet.htm` for more information. (accessed on August 1, 2017)

behind this rule was the reality that analysts had often recommended assets that the investment branch of their employer had to place in the market (although this action was not always to the benefit of the analysts' customers). However, as many commentators have pointed out, analysts still receive trading commissions entailing that the problem of non-aligned incentives between analysts and customers was only mitigated but not resolved. Trading commissions are, of course, an instrument for resolving a moral hazard problem between analysts and their employers. Consequently, they cannot be eliminated without creating an inefficiency at a different place. In other applications, it is not even possible to eliminate the misalignment of interests. In case of a political adviser, the bonus might simply be interpreted as a personal political preference. Such preferences and resulting differences in opinions appear to be inevitable. Nevertheless, the inefficiency established in this paper can be a rationale for regulations as the one mentioned above.

The paper explains why most advice relationships are short lived, which leads to the question: which advice relationships can last long? One possibility is that there is no conflict of interest, as mentioned in the previous paragraph. Another possibility is that players are very patient. The bounds T' in lemma 4 and $\bar{T}$ in theorem 2 converge to infinity as $\delta \to 1$. Intuitively, A is not tempted to get his bonus today quickly if there is almost no discounting and he might therefore be willing to give better advice today in order not to risk future bonus payments. Of course, the opposite also holds: For low δ, no meaningful advice is possible as the game with heavy discounting is similar to the static game where A will always recommend his bonus option.

6. Conclusion

This paper has analyzed the question why advisers are fired. Two reasons for their dismissal are identified in a repeated game model. First, incompetence; that is, advisers who are believed to be of low quality are fired. Second, (justified!) mistrust. Advisers who are believed to be competent are not afraid of being fired due to incompetence. In equilibrium, these advisers will therefore push their own agenda, i.e. recommend actions that foster their own benefit more than the decision maker's benefit. Consequently, the decision maker is indifferent to firing them and will do so with positive probability whenever he receives bad advice. The interplay of these two effects can imply that the decision maker receives the most informative recommendations from an adviser whose qualification is unclear. Such an adviser tries to give good advice because he is afraid to be perceived as being of so low quality that he is fired due to incompetence in case his advice turns out to be bad. The firing of competent advisers is inevitable in equilibrium but inefficient. Independent of qualification and beliefs, the expected length of the advice relationship is therefore limited although advice by a qualified adviser is efficient. Thus, the presence of private benefits for the adviser, like bonus payments, does not only lead to bad advice, it also implies that decision makers drop (eventually) the best advisers and end up with inferior (or no) advice.

The model of this paper helps to identify the effects mentioned above and paves the way for further research. For example, the literature on sell-side analysis in financial advice – see, for example, Fang and Yasuda (2009); Jackson (2005) – is concerned to what extent reputation effects can alleviate opportunistic behavior by analysts.[17] While this literature establishes empirically that more reputable analysts provide better predictions on average, this result is based on assuming either a binary or a linear functional form for this relationship. However, the model of this paper suggests a non-monotonic (possibly inversely U-shaped) relationship.[18] On the theoretical side, one might explore instruments that a decision maker could use to discipline advisers such as own (costly) acquisition of noisy information and consulting multiple advisers at the same point of time. Another interesting question to be raised is the effect of learning. Here, an adviser might receive more precise signals concerning the decision maker's needs as he interacts with the decision maker repeatedly. These and other possibilities are beyond the scope of the current paper and left for future research.

[17] "Sell-side analysis" refers to the situation where employees of a broker provide analysis and stock recommendations to potential customers for free in the hope of generating an order that yields a commission.

[18] The possibility of non-monotonicities seems to have escaped the attention of the authors, e.g. Fang and Yasuda (2009, p. 3736) write "Because analysts with a better reputation have greater long-term benefits to lose, theory predicts that they are more likely to refrain from opportunism." My paper has demonstrated that this argument, though plausible at first sight, might not be true in equilibrium.

Appendix

Proof of lemma 1: Recommending the bonus option yields an expected payoff of $1 + v^+/2 + v^-/2$ where v^+ (v^-) denotes the expected discounted future payoff stream after a fitting (non-fitting) recommendation. Recommending the non-bonus option has expected payoff $v^+/2 + v^-/2 < 1 + v^+/2 + v^-/2$. □

Proof of lemma 2: Suppose to the contrary that $\beta^+(\alpha^+)V(\alpha^+) < \beta^-(\alpha^-)V(\alpha^-)$. Then, A has an incentive to recommend options that do not fit DM's needs as this will give him the higher continuation value. Consequently, A will recommend either his bonus option or option l. Hence, a competent adviser will give (weakly) worse advice than an uniformed adviser which implies $\alpha^+ \leq \alpha^-$. In case of $\alpha^+ = \alpha^-$, $\beta^+(\alpha^+)V(\alpha^+) < \beta^-(\alpha^-)V(\alpha^-)$ cannot hold and in case $\alpha^+ < \alpha^-$ the equilibrium is not informative. □

Proof of lemma 3: The first item was explained in the main text. For the second item, let option l be the bonus option. Recommending option h will then give an expected payoff of $p\,\delta\beta^+(\alpha^+)V(\alpha^+) + \delta(1-p)\beta^-(\alpha^-)V(\alpha^-)$ and recommending option l will yield an expected payoff of $(1-p)\,\delta\beta^+(\alpha^+)V(\alpha^+) + \delta p\beta^-(\alpha^-)V(\alpha^-) + 1$. The former is higher than the latter if and only if (2) holds. □

Proof of proposition 1: DM's expected payoff from continuing is bounded from above by $\delta\alpha p/(1-\delta) + \delta(1-\alpha)/(2(1-\delta))$. For α sufficiently low (but strictly higher than 0), this upper bound is less than δW_O as $W_O > 1/(2(1-\delta))$ by (1). □

Proof of theorem 1: The first part is proven by contradiction, and the proof follows the argument in the main text. Suppose the statement was not true; i.e. suppose that there was a Markov equilibrium such that for no sequence $(\alpha_i)_{i=1}^{\infty}$ converging to 1 there exists an $\varepsilon > 0$ such that DM ends the game with at least ε probability at each element of the sequence.[19] This implies that for every $\varepsilon' > 0$ there exists an $\bar{\alpha}_{\varepsilon'} < 1$ such that DM continues with probability greater than $1 - \varepsilon'$ for all $\alpha \geq \bar{\alpha}_{\varepsilon'}$. Note that $\alpha - \alpha^-$ converges to zero as $\alpha \to 1$ (for any strategy A employs); see (5), which constitutes a lower bound on α^- for any strategy A might use. This implies the following: For every $T \in \mathbb{N}$ and $\varepsilon' > 0$ there is a $\alpha_{T\varepsilon'} \in (\bar{\alpha}_{\varepsilon'}, 1)$ such that DM's belief after T consecutive recommendations that did not fit DM's needs will still be above $\bar{\alpha}_{\varepsilon'}$. This implies that at the belief $\alpha_{T\varepsilon'}$ (for T high enough) A will find it strictly optimal to recommend his bonus option, even if it is option l: if the belief is $\alpha_{T\varepsilon'}$ and A observes that the bonus option is option l, then he can earn a deviation payoff of at least $1 + \delta(1-\varepsilon') + \delta^2(1-\varepsilon')^2 + \cdots + \delta^T(1-\varepsilon')^T$ by recommending the bonus option in this and the following T periods. Not recommending the bonus option would lead to a

[19]It is immaterial for this proof whether DM ends the game only when receiving a non-fitting recommendation or not. The reason is that A's signal is noisy. Hence, even when recommending option h there is a probability $1 - p > 0$ of not fitting DM's needs. If DM ends the game with probability $\tilde{\varepsilon} > 0$ if the recommendation does not fit, then he ends the game with at least probability $\varepsilon = (1-p)\tilde{\varepsilon}$.

payoff of at most $\delta/(1-\delta) = \delta + \delta^2 + \delta^3 + \dots$. For T large enough and ε' small enough, the deviation payoff is clearly higher than the upper bound on the payoff obtained by other strategies. By choosing T high and $\varepsilon > 0$ small enough, this establishes the claim that A always recommends the bonus option if his reputation is sufficiently high (i.e. $\alpha_{T\varepsilon'}$ or higher). But then $\alpha_{T\varepsilon'} = \alpha^+_{T\varepsilon'} = \alpha^-_{T\varepsilon'}$ and therefore $W(\alpha_{T\varepsilon'}) = 1/(2-2\delta) < W_O$ by (1); i.e. DM's best response is to end the game at belief $\alpha_{T\varepsilon'}$ contradicting the definition of $\bar{\alpha}_{\varepsilon'}$ (and $\alpha_{T\varepsilon'} > \bar{\alpha}_{\varepsilon'}$).

For the second part, note that piecewise continuity of the strategies implies piecewise continuity of the value functions. In particular, V is piecewise continuous and therefore has bounded total variation, which will be used below. Note that whenever A recommends the bonus option for sure (even if it is option l) at some α, then $\alpha = \alpha^+ = \alpha^-$ and therefore ending the game is DM's best response when α is reached, as continuing would lead to a payoff of $1/(2-2\delta) < W_O$.

Let $A_{\tilde{\varepsilon}}$ be the set of α such that (i) A recommends option h with positive probability even if it is not the bonus option, and (ii) $\beta^-(\alpha^-) > 1-\tilde{\varepsilon}$. If the last claim of the theorem holds, then $A_{\tilde{\varepsilon}} \cap (\bar{\alpha}, 1)$ is empty for $\tilde{\varepsilon} > 0$ small enough and $\bar{\alpha} < 1$ large enough. Suppose this is not the case, i.e. suppose $A_{\tilde{\varepsilon}} \cap (\bar{\alpha}, 1)$ is non-empty for all $\tilde{\varepsilon} > 0$ and $\bar{\alpha} < 1$. For $\alpha \in A_{\tilde{\varepsilon}}$, the difference $\beta^+(\alpha^+)V(\alpha^+) - \beta^-(\alpha^-)V(\alpha^-)$ is bounded from below by a strictly positive number by lemma 3 and the optimality of recommending option h if the bonus option is option l (i.e. the definition of $A_{\tilde{\varepsilon}}$). Together with $\beta^-(\alpha^-) > 1-\tilde{\varepsilon}$, this implies that $V(\alpha^+) - V(\alpha^-)$ is bounded from below as well for $\tilde{\varepsilon} > 0$ sufficiently small and $\alpha \in A_{\tilde{\varepsilon}}$.

Now note that $\alpha^+ - \alpha^-$ converges to zero as α approaches 1. Therefore, it is possible to construct an increasing sequence (α_i) of elements of $A_{\tilde{\varepsilon}}$ such that $\alpha^-_{i+1} \geq \alpha^+_i$. This can be done as for any given a^+_i there exists an $\hat{a}^+_i$ such that $\alpha^- > a^+_i$ for all $\alpha > \hat{a}^+_i$ (this is true because $\alpha - \alpha^-$ converges to zero as α converges to 1, and $A_{\tilde{\varepsilon}} \cap (\bar{\alpha}, 1)$ is non-empty for all $\bar{\alpha} < 1$). The construction of this sequence and the fact that $V(\alpha^+) - V(\alpha^-)$ is bounded from below by a strictly positive number then imply that V is a function of unbounded total variation. This, however, contradicts the piecewise continuity of V. Hence, $A_{\tilde{\varepsilon}} \cap (\bar{\alpha}, 1)$ has to be empty for $\bar{\alpha} < 1$ high enough and $\tilde{\varepsilon} > 0$ small enough. This establishes the last claim of the theorem.

The result that $W(\alpha) = W_O$ for almost all $\alpha > \bar{\alpha}$ follows directly from the fact that DM ends the game with positive probability for all $\alpha > \bar{\alpha}$ – the just proven last claim of the theorem – which is only optimal if $W(\alpha) = W_O$. □

Proof of lemma 4: Take an arbitrary subgame starting in an arbitrary period t and keep it fixed for the rest of the proof. First, I want to establish that there is at least one continuation path of play of length (up to) T' such that the probability that DM stops the game along this path is at least $\varepsilon' > 0$ for sufficiently high T' and sufficiently low $\varepsilon' > 0$. If DM's strategy is such that he stops with some probability in t, there is nothing to show. Hence, I will assume otherwise.

If T'' is sufficiently large and A's equilibrium strategy is such that he recommends his bonus option for the next T'' periods for sure (even if it is option l), then DM will find it optimal to end the game immediately as his equilibrium payoff from continuing would be bounded from above by $(1+\delta+\cdots+\delta^{T''})/2+\delta^{T''+1}p/(1-\delta)$ which is less than W_O for T'' high enough by (1). Hence, by choosing T'' sufficiently high there will be some continuation path on which A recommends option h (with strictly positive probability) in period $t' \in \{t, \dots, T''\}$ even if the bonus option is option l. Note that A will do so only if there is a continuation path (from t') onward on which DM ends the game with some probability $\varepsilon'' > 0$ within the next T''' periods (for some T''' high enough): Otherwise, A could profitably deviate by recommending his bonus option for the next T''' periods. Letting $T' = T'' + T'''$ and $\varepsilon' = \varepsilon''$ yields the result that starting from t there is at least one continuation path of play of length (up to) T' such that the probability that DM ends the game along this path is at least $\varepsilon' > 0$.

Next, I will show that $\varepsilon' = (1-\delta)/2$ and $T' = \lceil 2\log(1-\delta)/\log(\delta) - 1\rceil$ have the desired properties. Suppose otherwise, i.e. suppose there is no continuation path of of length T' such that DM ends the game on this path with probability ε' or higher. I will show that in this case A will have an incentive to deviate if the bonus option is option l. By recommending the bonus option, A can achieve a payoff of at least $1+(1-\varepsilon')(\delta+\delta^2+\cdots+\delta^{T'})$. By sticking to his equilibrium strategy (i.e. not recommending the bonus option) A will achieve a payoff of at most $0+\delta/(1-\delta)$. With ε' and T' chosen as above, however, the lower bound on the deviation profit is higher than the upper bound of the equilibrium profit, i.e. the deviation is profitable:

$$\begin{aligned}
T' = \left\lceil 2\frac{\log(1-\delta)}{\log(\delta)} - 1\right\rceil \quad &\geq \quad 2\frac{\log(1-\delta)}{\log(\delta)} - 1 \\
\Leftrightarrow (T'+1)\log(\delta) \quad &\leq \quad 2\log(1-\delta) \\
\Leftrightarrow \delta^{T'+1} \quad &\leq \quad (1-\delta)^2 \\
\Leftrightarrow \frac{\delta^{T'+1}}{1-\delta} \quad &\leq \quad 1-\delta \\
\Leftrightarrow \delta\left(1+\delta^{T'+1}+\delta^{T'+2}+\dots\right) \quad &\leq \quad 1 \\
\Leftrightarrow (1-\delta)\left(\delta+\delta^2+\cdots+\delta^{T'}\right)+\delta^{T'+1}+\delta^{T'+2}+\dots \quad &\leq \quad 1 \\
\Leftrightarrow \delta+\delta^2+\dots \quad &\leq \quad 1+\delta\left(\delta+\delta^2+\cdots+\delta^{T'}\right) \\
\Leftrightarrow \frac{\delta}{1-\delta} \quad &\leq \quad 1+(1-2\varepsilon')\left(\delta+\delta^2+\cdots+\delta^{T'}\right)
\end{aligned}$$

which implies that the lower bound of the deviation payoff (which is the right hand side but without multiplying ε' by 2) is strictly higher than the upper bound for the equilibrium payoff. This establishes that ε' and T' have the desired property.

As the continuation path on which DM's probability to end the game is (at least) ε' has positive probability under equilibrium play, by the assumption that A is uncertain (i.e. $p < 1$), it follows that the game ends with probability $\gamma\varepsilon' > 0$ in the next T' periods where γ is a lower bound on the probability of the path occurring under equilibrium play which can be chosen independent of the specifics of the equilibrium and the belief, i.e. depending only on the precision of A's signal p. For example, $\gamma = (1-p)^{T'}$ works and will be used in the remainder.

Hence, the probability that DM does not end the game within $2T'$ periods is at most $(1-\gamma\varepsilon')^2$. Iterating yields that the probability that DM does not end the game within mT' periods is at most $(1-\gamma\varepsilon')^m$. Let m' be such that $\varepsilon > (1-\gamma\varepsilon')^{m'}$, and let $T_\varepsilon > m'T'$. Using $\gamma = (1-p)^{T'}$ and T', ε' as derived above yields

$$T_\varepsilon = \left\lceil \frac{\log(\varepsilon)}{\log\left(1-(1-p)\varepsilon'\right)} \right\rceil T'.$$

The result follows. □

Proof of theorem 2: Lemma 4 states that the probability that the game lasts longer than T_ε periods is *at most* ε. As I want to derive an upper bound on the expected length, I can assume that the probability that the game lasts longer than T_ε periods is *exactly* ε. As it simplifies the derivation and since I am only interested in an upper bound, I will actually assume that the probability that the game lasts longer than

$$\tilde{T}_\varepsilon = \frac{\log(\varepsilon)}{\log\left(1-(1-p)^{T'}\varepsilon'\right)} T' + T'$$

equals ε for $\tilde{T}_\varepsilon > T'$, which again will increase the expectation as $\tilde{T}_\varepsilon \geq T_\varepsilon$. That is, I assume that the game lasts at least T' periods (which again increases the expectation). Rearranging yields that the probability that the game's length is $\hat{T} > T'$ or less is $1 - e^{(\hat{T}-T')/B}$ where $B = T'/\log\left(1-(1-p)^{T'}\varepsilon'\right)$. Note that $B < 0$. The corresponding density is $-e^{(\hat{T}-T')/B}/B$. This allows to compute an upper bound on the expected length of the game as

$$\begin{aligned} T' + \int_{T'}^{\infty} -\frac{\hat{T}e^{(\hat{T}-T')/B}}{B}\, \mathrm{d}\hat{T} &= T' + \left[-\hat{T}e^{(\hat{T}-T')/B} + Be^{(\hat{T}-T')/B}\right]_{T'}^{\infty} \\ &= 2T' - B = T'\left(2 - \frac{1}{\log(1-(1-p)^{T'}\varepsilon')}\right). \end{aligned}$$

□

Proof of proposition 2:

Sufficiency of (6): The proof is by construction. Let (7) be satisfied and define

$$\begin{aligned}
\underline{\alpha} &= \frac{W_O(1-\delta)-1/2}{p-1/2}, \\
\bar{V} &= \frac{4p-1}{4p-2}.
\end{aligned}$$

Note that $\underline{\alpha} < 1$ by (1). Furthermore, $\bar{V} > 1$ as $p > 3/4$ by (6).

I will show that the following value functions and strategies constitute a Markov equilibrium:

$$\begin{aligned}
V(\alpha) &= \begin{cases} 1 & \text{if } \alpha \leq \underline{\alpha} \\ \bar{V} & \text{if } \alpha > \underline{\alpha} \end{cases} \\
W(\alpha) &= \begin{cases} 1/2 + \delta W_O & \text{if } \alpha \leq \underline{\alpha} \\ W_O & \text{if } \alpha > \underline{\alpha} \end{cases} \\
s(\alpha, h) &= 1 \\
s(\alpha, l) &= \begin{cases} 0 & \text{if } \alpha \leq \underline{\alpha} \\ \frac{(1-\delta)W_O - 1/2}{\alpha(p-1/2)} & \text{if } \alpha > \underline{\alpha} \end{cases} \\
\beta^-(\alpha^-) &= 0 \\
\beta^+(\alpha^+) &= \begin{cases} 0 & \text{if } \alpha \leq \underline{\alpha} \\ \frac{1}{\bar{V}\delta(2p-1)} & \text{if } \alpha > \underline{\alpha} \end{cases}
\end{aligned}$$

where $s(\alpha, b)$ gives the probability with which A recommends option h if option b is his bonus option. Note that $0 \leq 1/(\bar{V}\delta(2p-1)) \leq 1$ by (6).

The above is a Markov equilibrium if and only if the strategies are mutual best responses and the value functions are consistent with the strategies; i.e. if and only if the Bellman equations

$$\begin{aligned}
V(\alpha) &= \max_{s_h, s_l} \left\{ \frac{1}{2} \left[s_h + \delta q^+(s_h)\beta^+(\alpha^+)V(\alpha^+) + \delta q^-(s_h)\beta^-(\alpha^-)V(\alpha^-) \right] \right. \\
&\quad \left. + \frac{1}{2} \left[(1-s_l) + \delta q^+(s_l)\beta^+(\alpha^+)V(\alpha^+) + \delta q^-(s_l)\beta^-(\alpha^-)V(\alpha^-) \right] \right\} \\
W(\alpha) &= \max_{\beta^+, \beta^-} \left\{ q^+(\alpha) + \delta q^+(\alpha)\beta^+ W(\alpha^+) + \delta q^-(\alpha)\beta^- W(\alpha^-) \right. \\
&\quad \left. + \delta(1 - q^+(\alpha)\beta^+ - q^-(\alpha)\beta^-)W_O \right\}
\end{aligned}$$

hold and the strategies are the maximizing arguments (where $q^+(s_i) = s_i p + (1-s_i)(1-p)$ is the probability of giving a fitting recommendation and $q^-(s_i) = 1 - q^+(s_i)$ is the counter-probability; similarly and with a slight abuse of notation, $q^+(\alpha) = (1-\alpha)/2 + \alpha[s(\alpha,h)p + (1-s(\alpha,h))(1-p)]/2 + \alpha[s(\alpha,l)p + (1-s(\alpha,l))(1-p)]/2$ is the expected probability of getting a fitting recommendation and $q^-(\alpha) = 1 - q^+(\alpha)$). Consistency of the value functions with the strategies is straightforward to verify, and I will therefore only briefly explain why the strategies are indeed the maximizing arguments.

Recall that A wants to recommend option h in case the bonus option is option l if and only if $\beta^+(\alpha^+)V(\alpha^+) - \beta^-(\alpha^-)V(\alpha^-) \geq 1/(\delta(2p-1))$, see lemma 3. DM's strategy is chosen such that A is indifferent for $\alpha > \underline{\alpha}$. For $\alpha \leq \underline{\alpha}$, recommending the bonus option is clearly optimal as $\beta^+(\alpha^+) = \beta^-(\alpha^-) = 0$.

DM acts optimally when $\alpha \leq \underline{\alpha}$ as in this case $q^+(\alpha) = 1/2 = q^-(\alpha)$ and $\alpha^+ = \alpha^- = \alpha$ and therefore stopping the game gives W_O (in the next period) while continuing gives $1/2 + \delta W_O$ which is less than W_O by (1). For $\alpha > \underline{\alpha}$, $s(\alpha, l)$ is chosen such that DM is indifferent between continuing and stopping after a fitting recommendation (note that $q^+(\alpha) = (1-\delta)W_O$). After a non-fitting recommendation, DM is either indifferent (if $\alpha^- > \underline{\alpha}$) or strictly prefers stopping (if $\alpha^- \leq \underline{\alpha}$).

Necessity of (6): This proof is split up into two lemmas:

Lemma 5. *In every communication equilibrium,* $\lim_{\alpha\to 1} s(\alpha, l) > 0$.

Proof of lemma 5: Two preliminary observations: First, if A is willing to recommend option h at belief α even though the bonus option is option l, then necessarily $V(\alpha^+) > 1/\delta$ (as A could get 1 for sure in this period by recommending the bonus option). Second, if $V(\alpha) > 1$, then DM has to play continue (after a fitting recommendation) with positive probability as otherwise $V(\alpha) \leq 1$.

Now let $\mathcal{A} = \{\alpha : \ s(\alpha, l) > 0\}$ and $\bar{a} = \sup \mathcal{A}$. Note that $\mathcal{A}$ is the union of a finite number of intervals by the assumption that the equilibrium strategies are piecewise continuous. If $\bar{a}$ was below 1, then $s(\alpha, l) = 0$ for $\alpha > \bar{a}$ which would imply that DM's best response at these α is to always end the game and therefore $V(\alpha) = 1$ and $W(\alpha) = W_O$ for all $\alpha > \bar{a}$.

I want to show that $\bar{a} = 1$. Note that $\bar{a} = 1$ holds if $\lim_{\alpha\nearrow\bar{a}} s(\alpha, l) > 0$: As $\lim_{\alpha\nearrow\bar{a}} s(\alpha, l) > 0$ implies that $s(\alpha, l) > \varepsilon'$ on $(\bar{a} - \varepsilon, \bar{a})$ for some ε, $\varepsilon' > 0$ by the assumption that s is piecewise continuous, it follows that $\alpha^+ > \bar{a}$ for $\alpha < \bar{a}$ sufficiently close to $\bar{a}$. Following the first preliminary observation above this requires $V(\alpha^+) > 1/\delta$ which contradicts that $V(\alpha) = 1$ for all $\alpha > \bar{a}$. Note that this argument shows more generally that $\alpha^+ \leq \bar{a}$ for all $\alpha \in \mathcal{A}$ (with strict inequality if $\bar{a} < 1$ or $\alpha < \bar{a}$).

Consequently, I only have to rule out $\lim_{\alpha\nearrow\bar{a}} s(\alpha, l) = 0$ to show that $\bar{a} = 1$. To do so, I will show that DM wants to end the game for some $\alpha < \bar{a}$ arbitrarily close to $\bar{a}$ if $\lim_{\alpha\nearrow\bar{a}} s(\alpha, l) = 0$ which contradicts that $s(\alpha, l) > 0$ is a best response for these beliefs. Suppose $\lim_{\alpha\nearrow\bar{a}} s(\alpha, l) = 0$. As the equilibrium is assumed to be piecewise continuous, $s(\alpha, l)$ is continuous on $(\bar{a} - \varepsilon, \bar{a})$ for some $\varepsilon > 0$. Hence, for every $\varepsilon'' > 0$, there exists an $\varepsilon' > 0$ such that $s(\alpha, l) < \varepsilon''$ for $\alpha \in (\bar{a} - \varepsilon', \bar{a})$. Note furthermore that $\lim_{\alpha\nearrow\bar{a}} \alpha - \alpha^- = 0$ as competent and incompetent type use (in the limit) the same strategy as $\alpha \nearrow \bar{a}$. Consequently, for every $T \in \mathbb{N}$ and every $\varepsilon' > 0$, there exists an $\alpha_{T\varepsilon'} \in (\bar{a} - \varepsilon', \bar{a})$ such that, starting from belief $\alpha_{T\varepsilon'}$, the updated belief after T consecutive recommendation that did not fit DM's needs will still be above $\bar{a} - \varepsilon'$. Starting from belief $\alpha_{T\varepsilon'}$, the updated belief of the next T periods will therefore be in $(\bar{a} - \varepsilon', \bar{a})$ as the previous paragraph established that $\alpha^+ \leq \bar{a}$ for all $\alpha \in \mathcal{A}$. Choosing ε' such that (i) $s(\alpha, l) < \varepsilon''$ for $\alpha \in (\bar{a} - \varepsilon', \bar{a})$ and (ii) $(\bar{a} - \varepsilon', \bar{a}) \subset \mathcal{A}$ yields the following upper bound on $W(\alpha_{T\varepsilon'})$ (recall that by the definition of $\mathcal{A}$ continuing is a best response for all $\alpha \in \mathcal{A}$):

$$\delta^{T+1}\frac{p}{1-\delta} + \sum_{t=0}^{T} \delta^t \left[(1-\alpha)\frac{1}{2} + \alpha\left(p(\frac{1}{2} + \varepsilon''\frac{1}{2}) + (1-p)\frac{1}{2}\right)\right].$$

This upper bound is strictly less than the W_O by assumption (1) for ε'' sufficiently small and T sufficiently high. Hence, ending the game is DM's unique best response at $\alpha_{T\varepsilon'}$ for T high enough and $\varepsilon' > 0$ small enough which contradicts that $\alpha_{T\varepsilon'} \in \mathcal{A}$. □

Lemma 6. *(7) holds in every communication equilibrium.*

Proof of lemma 6: By piecewise continuity of the equilibrium strategies, there exists an $\bar{a} < 1$ such that V is continuous for all $\alpha \in (\bar{a}, 1)$. By lemma 5 and piecewise continuity, $\bar{a}$ can be chosen high enough to ensure $s(\alpha, l) > 0$ for $\alpha \in (\bar{a}, 1)$. By theorem 1, $\bar{a}$ can be chosen high enough to also ensure $\min\{\beta^-(\alpha^-), \beta^+(\alpha^+)\} < 1$ for $\alpha \in (\bar{a}, 1)$. Note that firing the adviser, i.e. $\beta^-(\alpha^-) < 1$ or $\beta^+(\alpha^+) < 1$, is only a best response for high α if $s(\alpha, l) < 1$ (by assumption (1)). Hence, there exists an $\bar{a} < 1$ such that (i) $0 < s(\alpha, l) < 1$ for all $\alpha \in (\bar{a}, 1)$, (ii) V is continuous on $\alpha \in (\bar{a}, 1)$ and (iii) β^+ and β^- are continuous on $\alpha \in (\bar{a}, 1)$. In the remainder of this proof, only beliefs in $\alpha \in (\bar{a}, 1)$ are considered, i.e. (i), (ii) and (iii) are assumed to hold, and the qualifier "for $\alpha \in (\bar{a}, 1)$" is omitted.

As A uses a mixed strategy, he has to be indifferent between recommending either of the two options if option l is the bonus option. That is, the indifference condition $\beta^+(\alpha^+)V(\alpha^+) - \beta^-(\alpha^-)V(\alpha^-) = 1/(\delta(2p-1))$ holds; see lemma 3.

As A is indifferent in case option l is the bonus option (and always recommends the bonus option if it is option h), one way to achieve his expected value $V(\alpha)$ is to always recommend his bonus option. Since the bonus option fits DM's needs with probability 1/2, this gives

$$V(\alpha) = 1 + \delta\left(\frac{1}{2}\beta^+(\alpha^+)V(\alpha^+) + \frac{1}{2}\beta^-(\alpha^-)V(\alpha^-)\right). \tag{8}$$

Using the indifference condition to eliminate $\beta^+(\alpha^+)V(\alpha^+)$ gives

$$V(\alpha) = 1 + \frac{1}{2(2p-1)} + \delta\beta^-(\alpha^-)V(\alpha^-). \tag{9}$$

V is continuous and bounded; e.g. bounded from above by $1/(1-\delta)$, and bounded from below by 0. Therefore, $\tilde{V} \equiv \lim_{\alpha\to 1} V(\alpha)$ exists. Furthermore, $\lim_{\alpha\to 1} \alpha^- = \alpha$ because updating stops as the belief is approaching 1. The continuity of V implies therefore that $\lim_{\alpha\to 1} V(\alpha) = \lim_{\alpha\to 1} V(\alpha^-)$. Strategies are continuous by assumption and therefore $\lim_{\alpha\to 1} \beta^-(\alpha^-)$ exists and will – with a slight abuse of notation – be denoted as $\beta^-(1)$. Taking limits on both sides of (9) yields (after rearranging)

$$\tilde{V} = \frac{4p-1}{(4p-2)(1-\delta\beta^-(1))}.$$

As $\beta^-(1)$ is a probability and therefore in $[0,1]$, the previous equation yields

$$\tilde{V} \in \left[\frac{4p-1}{4p-2}, \frac{4p-1}{(4p-2)(1-\delta)}\right]. \tag{10}$$

In equation (8), one could also use the indifference condition to eliminate $\beta^-(\alpha^-)V(\alpha^-)$ (instead of $\beta^+(\alpha^+)V(\alpha^+)$). This yields

$$V(\alpha) = 1 - \frac{1}{2(2p-1)} + \delta\beta^+(\alpha^+)V(\alpha^+).$$

Taking again limits on both sides results (after rearranging) in

$$\tilde{V} = \frac{4p-3}{(4p-2)(1-\delta\beta^+(1))}.$$

As $\beta^+(1)$ is a probability and therefore in $[0,1]$, this implies that

$$\tilde{V} \in \left[\frac{4p-3}{4p-2}, \frac{4p-3}{(4p-2)(1-\delta)}\right]. \tag{11}$$

Taking (10) and (11) together yields the result. □

Finally, note that the upper bound of the interval in (7) is higher than the lower bound of the interval if and only if (6) holds. That is, if (6) does not hold, $\tilde{V}$ cannot both satisfy (10) and (11) at the same time and a communication equilibrium cannot exist. □

Proof of corollary 1: By (7), $\lim_{\alpha\to 1} \leq (4p-3)/((4p-2)(1-\delta))$ in every such equilibrium. Rearranging

$$\frac{4p-3}{(4p-2)(1-\delta)} < \frac{1}{2(1-\delta)}$$

yields – given that $p > 1/2$ – $p < 1$, which is true by assumption. □

Proof of proposition 3: The sum of expected payoffs in a given period has to be less than $1/2 * 1 + 1/2 * 2$ as with probability 1/2 the bonus option coincides with the option fitting DM's needs (allowing for a maximal possible payoff of 2) and with probability 1/2 the two do not coincide which means that only one of the two players can have a payoff of 1. Consequently, the sum of expected continuation payoffs in a given subgame has to be less than $(3/2)/(1-\delta)$.

DM can ensure himself an expected payoff of at least $1/2 + \delta W_O$ in any subgame by stopping the game at the end of the current period without transferring any wage payments to A. (1/2 is the payoff DM expects to receive from an incompetent type in a given period, and by informativeness of the equilibrium, 1/2 is also the minimum expected payoff from a competent type.) The upper bound on continuation welfare and the lower bound on DM's continuation payoff yield an upper bound on A's expected continuation payoff equal to

$$\bar{V} = (3/2)/(1-\delta) - 1/2 - \delta W_O = \frac{1-\eta\delta}{1-\delta} \tag{12}$$

where $\eta = (1-\delta)W_O - 1/2$ and $\eta > 0$ by (1).

Now suppose – contrary to the proposition – that there exists an $\bar{\alpha} < 1$ such that DM continues with probability of at least $1-\varepsilon$ whenever his belief is above $\bar{\alpha}$. As belief updating becomes arbitrarily slow when α is close to 1, there exists an α_T for every $T \in \mathbb{N}$ such that an adviser of reputation α_T (or higher) will still have a reputation above $\bar{\alpha}$ after T periods, i.e. even if he recommends for T consecutive times a non-fitting option. An adviser with reputation of (at least) α_T can guarantee himself a continuation payoff of

$$\hat{V} = \sum_{t=0}^{T} (\delta(1-\varepsilon))^t = \frac{1-\delta^{T+1}}{1-\delta(1-\varepsilon)}$$

by recommending his bonus option for the next T periods. For T sufficiently high, however, $\hat{V}$ is above the upper bound for A's equilibrium continuation payoffs $\bar{V}$:

$$\begin{aligned}
\frac{1-\delta^{T+1}}{1-\delta(1-\varepsilon)} &> \frac{1-\eta\delta}{1-\delta} \\
\Leftrightarrow 1-\delta-(1-\delta)\delta^{T+1} &> 1-\eta\delta-\delta(1-\varepsilon)+\eta\delta^2(1-\varepsilon) \\
\Leftrightarrow (1-\varepsilon)(1-\eta\delta) &> 1-\eta+(1-\delta)\delta^T \\
\Leftrightarrow \frac{(\eta-\delta^T)(1-\delta)}{1-\eta\delta} &> \varepsilon
\end{aligned}$$

which is true by assumption for sufficiently high T and, for example, $\varepsilon = \eta(1-\delta)/(2-2\eta\delta) > 0$. $\hat{V} > \bar{V}$ contradicts that the supposed equilibrium exists and therefore establishes the desired result. $\square$

Proof of proposition 4: This proof will use the upper bound $\bar{V}$ on A's continuation payoff derived in the proof of proposition 3, see equation (12).

First, it is shown by contradiction that – starting from an arbitrary subgame – DM stops the game with at least probability ε' along some path of length T'. Suppose this was not the case. A can then guarantee himself an expected payoff of

$$\underline{V} = 1 + \frac{\delta(1-\varepsilon')-\delta^{T'+1}}{1-\delta(1-\varepsilon)}$$

by recommending the bonus option for the next T' periods. (This – very conservative – lower bound assumes a worst case where DM stops the game with probability ε' in each of the following T' periods and stops the game with probability 1 in period $T'+1$.) The contradiction emerges as $\underline{V} > \bar{V}$:

$$\begin{aligned}
1+\frac{\delta(1-\varepsilon')-\delta^{T'+1}}{1-\delta(1-\varepsilon)} &> \frac{1-\eta\delta}{1-\delta} \\
\Leftrightarrow 1-\delta-(1-\delta)\delta^{T'+1} &> 1-\delta\eta-\delta(1-\varepsilon')+\delta^2(1-\varepsilon')\eta \\
\Leftrightarrow (1-\delta\eta)(1-\varepsilon')-(1-\eta) &> (1-\delta)\delta^{T'} \\
\Leftrightarrow T' &> \frac{log\left(\frac{(1-\delta\eta)(1-\varepsilon')-1+\eta}{1-\delta}\right)}{log(\delta)}
\end{aligned}$$

which holds true given the values of T' and ε' stated in the proposition. Hence, the game ends with at least ε' probability along some path of length T' in equilibrium (starting from any subgame). This implies that the probability that DM continues for mT' periods is at most $(1-\gamma\varepsilon')^m$ where $\gamma = (1-p)^{T'}$ is a lower bound on the probability of the specific path of recommendations having at least ε' probability of ending the game. Now take an arbitrary $\varepsilon > 0$ and let $m \in \mathbb{N}$ be sufficiently high such that $(1-\gamma\varepsilon')^m \leq \varepsilon$, e.g. $m = \lceil log(\varepsilon)/log(1-\gamma\varepsilon')\rceil$. Then, the probability that the game continues for T_ε or more periods is less than ε. This establishes the first result.

The second result follows from the first result in exactly the same way in which theorem 2 followed from lemma 4, i.e. the proof of theorem 2 applies with T_ε, ε' and T' as stated in proposition 4. □

Proof of proposition 5: Let

$$W_O = \frac{\alpha_0(p - 1/2) + 1}{2(1 - \delta)}$$

and $\underline{\alpha} = \alpha_0/2$. Consider the equilibrium strategies and value functions as in the sufficiency part of the proof of proposition 2 (where V and s denote now the strategies for each adviser). As shown there, these strategies are an equilibrium given W_O. As $W(\alpha_0) = W_O$, W_O is indeed DM's value of resetting. □

References

Aghion, P. and M. Jackson (2016). Inducing leaders to take risky decisions: dismissal, tenure, and term limits. *American Economic Journal: Microeconomics 8*(3), 1–38.

Aumann, R. J. and S. Hart (2003). Long cheap talk. *Econometrica 71*(6), 1619–1660.

Benabou, R. and G. Laroque (1992). Using privileged information to manipulate markets: Insiders, gurus, and credibility. *Quarterly Journal of Economics 107*(3), 921–958.

Brandenburger, A. and B. Polak (1996). When managers cover their posteriors: Making the decisions the market wants to see. *RAND Journal of Economics 27*(3), 523–541.

Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica 50*(6), 1431–1451.

Cripps, M. W., G. J. Mailath, and L. Samuelson (2004). Imperfect monitoring and impermanent reputations. *Econometrica 72*(2), 407–432.

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

Fama, E. F. (1980). Agency problems and the theory of the firm. *Journal of Political Economy 88*(2), 288–307.

Fang, L. and A. Yasuda (2009). The effectiveness of reputation as a disciplinary mechanism in sell-side research. *Review of Financial Studies 22*(9), 3735–3777.

Golosov, M., V. Skreta, A. Tsyvinski, and A. Wilson (2014). Dynamic strategic information transmission. *Journal of Economic Theory 151*, 304–341.

Holmström, B. (1982). Managerial incentive problems: A dynamic perspective. Reprinted in *Review of Economic Studies* (1999) *66*(1), 169–182.

Jackson, A. R. (2005). Trade generation, reputation, and sell-side analysts. *Journal of Finance 60*(2), 673–717.

Jullien, B. and I.-U. Park (2014). New, like new, or very good? reputation and credibility. *Review of Economic Studies 81*(4), 1543–1574.

Klein, N. and T. Mylovanov (2016). Will truth out? – an advisor's quest to appear competent. mimeo, University of Pennsylvania.

Krishna, V. and J. Morgan (2008). Cheap talk. In S. N. Durlauf and L. E. Blume (Eds.), *New Palgrave Dictionary of Economics*. Basingstoke: Palgrave Macmillan.

Li, W. (2007). Changing one's mind when the facts change: incentives of experts and the design of reporting protocols. *Review of Economic Studies 74*(4), 1175–1194.

Mailath, G. J. and L. Samuelson (2001). Who wants a good reputation? *Review of Economic Studies 68*(2), 415–441.

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

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

Ottaviani, M. and P. N. Sørensen (2006b). Reputational cheap talk. *RAND Journal of Economics 37*(1), 155–175.

Park, I.-U. (2005). Cheap talk referrals of differentiated experts in repeated relationship. *RAND Journal of Economics 36*(2), 391–411.

Prendergast, C. and L. Stole (1996). Impetuous youngsters and jaded old-timers: Acquiring a reputation for learning. *Journal of Political Economy 104*(6), 1105–1134.

Renault, J., E. Solan, and N. Vieille (2013). Dynamic sender–receiver games. *Journal of Economic Theory 148*(2), 502–534.

Schottmüller, C. (2016). Too good to be truthful: Why competent advisers are fired. TILEC Discussion Paper No. 2016-007.

Sobel, J. (1985). A theory of credibility. *Review of Economic Studies 52*(4), 557–573.

Wiseman, T. (2008). Reputation and impermanent types. *Games and Economic Behavior 62*(1), 190–210.