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

Christoph Schottmüller†
University of Copenhagen and Tilec

August 10, 2016

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 is more attractive for him to push his own agenda because he is less likely to be fired for incompetence. Consequently, competent advisers are also fired with positive probability. Firing is least likely if the decision maker is unsure about the adviser's type.

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 not surprising that advice by specialized experts is important in 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 personal and search engine operators) can get bonuses if their customers buy specific products, 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

*I want to thank Ole Jann, Jesper Rüdiger, Peter Norman Sørensen, Jörgen Weibull and audiences at the University of Copenhagen and SING12 in Odense for helpful comments.

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

lobbying efforts by its producer. Even ex post it is hard to detect whether these worries were justified because advice in all these areas is complex and even the best possible advice could turn out wrongly once in a while.

Another common feature of these examples is the repeated nature of advice: Most people tend to get advice from the same adviser several times and switch advisers only from times to times. 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 who are viewed as very competent are also fired occasionally: In 2003, financial analyst Jack Grubman was banned by the Security and Exchange Commission from the financial industry for life and fined fifteen million dollar for misconduct. Grubman had used his good reputation to pursue his own goals instead of his customers' ones when he gave a buy recommendation for AT&T as part of a plan to get his kids into 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 a year later as consultant, the company's stock price increased. The problem was that Grubman apparently (ab-)used his good reputation by misrepresenting his information and thereby manipulating his followers for his own personal benefit.

As an example from political advice, consider the firing of Roger Stone as Donald Trump's campaing 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 the firing 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." That is, the firing was not due to incompetence but due 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 dismissed or even killed their most prominent adivsers when these advisers were too competent and perceived as a threat to the throne.[2]

In all these examples, a very competent adviser is mistrusted and fired after committing some "mistake" that made the decision maker doubt whether the adviser acted

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

[2] Famous in this sense is the Ottoman Sultan Suleiman the Magnificient who killed not only his Grand Vizier and childhood friend Pargali Ibrahim Pasha (after Ibrahim committed the mistake of using the title "serasker sultan") but also his own son and designated heir Mustafa for this reason.

in the decision maker's best interest or whether he was (ab)using his power to push his own agenda instead. This paper argues that these situations are typical. More specifically, advisers are fired not *although* they had reputation for being competent but *because* they had this reputation for competence. That is, they might have kept their positions – and possibly not even committed the same 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 doubtful is facing the danger of being dismissed because of incompetence if his advice turns out to be bad (which strengthens the decision maker's initial doubts). Consequently, the advisor has strong incentives to act in the decision maker's best interest in order to keep his position. An adviser who is believed – with high probability – to be competent, however, has more freedom because the danger of being fired *due to incompetence* in the near future is negligible for him: 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. 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.[3]

Figure 1 shows the reasons for firing an adviser who is in fact competent at 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, 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 values an adviser most when he is uncertain about the quality of the adviser because this uncertainty will incentivize the adviser to give good advice. The decision maker does not value the adviser for very low or very high beliefs either because of likely incompetence or because of moral hazard.

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 model is a repeated game in which the adviser recommends one of finitely many

[3] To be more precise, decision maker and adviser might in equilibrium use mixed strategies when the adviser is believed to be competent: The decision maker will fire the adviser with some probability when a recommendation turns out to be bad. This gives the adviser some incentives to give not too bad advice in order to avoid being fired. However, the quality of advice will in equilibrium 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.

options to 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 usually 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 has 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 fitted his needs only after he 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 (or a finitely repeated game) because the adviser would always recommend his bonus option if he did not face the threat of losing future bonus payments. 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, the adviser is also fired with positive probability whenever he recommends an option not fitting 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. While the result that the adviser is fired (with some probability) if the belief in competence is high is naturally stated in terms of Markov equilibrium, the results on the expected length of the game are more general, i.e. they also hold in non-Markovian 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, 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).[4] The exact structure of the payoffs

[4]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 adviser 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

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 pushes therefore always for more leftist policies than the decision maker would like. In my setup, the bias is not in a certain direction but 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. [5] 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 on advice to find out which of several repair shops is specialized in fixing the problem at hand. In contrast to the current paper, the adviser knows in these papers the state of the world perfectly and reputation for competence does consequently 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 to face 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. 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), analyze how an expert who wants to maximizes his reputation for being competent will misrepresent his information. The resulting incentives to misrepresent towards the prior are not present in the current paper as the decision maker will have a uniform prior over a discrete set which makes it impossible to misrepresent towards the prior. Furthermore, the expert wants to maximize his expected bonus stream (not his reputation per se) which leads to the aforementioned moral hazard problem that drives the results of my paper. Papers where misreporting of information is driven by the interplay of (prior) beliefs and reputation concerns in a repeated game

advice or not.

[5]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).

setup include Prendergast and Stole (1996), Morris (2001), Ely and Välimäki (2003), Li (2007) and Klein and Mylovanov (2016).

The outline of the paper is as follows. The next section introduces the model and describes the solution concept 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 provides an illustrative numerical example and points out a commitment problem on the side of the adviser. Section 4 discusses the results and their implications and also shows that it is possible to relax some of the assumptions of the model. Section 5 concludes. The first part of the appendix outlines an algorithm for finding equilibria numerically; the second part contains the proofs of the results in the main text.

2. Model

The model is an infinitely repeated game. At the end of each period the decision maker (DM) can decide to end the game and consume his outside option W_O. If DM did not end the game in previous periods, the stage game in period t is the following. The adviser (A) receives a signal about DM's needs in this period and recommends one of n options to DM. DM will get a payoff of 1 (in this period) if the recommendation indeed fits his needs, otherwise DM receives a payoff of 0. Exactly one of the n options fits DM's needs and whether the recommendation fitted DM's needs or not is observed by both A and DM. One of the n options carries a bonus for A. If A recommends the bonus option, he receives a payoff of 1 in this period. Otherwise, A receives a payoff of 0. The identity of the bonus option is privately observed by A. Ex ante, each option is equally likely to be the bonus option.

DM considers – ex ante, i.e. before getting the advice – all options to be equally likely to fit his needs, i.e. his prior is that each option fits his needs with probability $1/n$. Furthermore, each option is – from DM's point of view – equally likely to be the one giving A the bonus. Consequently, DM cannot infer from the identity of the recommended option how likely it is that A recommended the bonus option.[6] The model considers the case where options and needs are not related across periods, e.g. DM asks for advice on a different question each period and the n options are actually different ones.

A has one of two types. A's type is time invariant and privately known by A. If

[6]One interpretation is that DM knows that there are n options but he cannot distinguish them, e.g. he is unaware of what the possible options are and only learns that a specific option exists if it is recommended to him (this is 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. Also in financial advice, customers usually do not know the bonuses that are associated with all possible investments. In these cases, advice is an experience good.

A is *competent*, he receives an informative (though not perfectly informative) signal which allows him to order the n options according to their likelihood. For notational simplicity, I will refer to the $k-$most likely option as option k (but keep in mind that only a competent A is aware of which option is the k-most likely while DM is not!). The posterior of a competent adviser is denoted by $(p_1, p_2, \ldots, p_n)$ where $p_k > p_{k+1}$ is the posterior probability that A assigns to option k being the one fitting DM's needs. Clearly, $\sum_{k=1}^{n} p_k = 1$. It is assumed that $p_1 < 1$ and $p_n > 0$, i.e. A cannot completely rule out any option.[7] All p_i are assumed to be time invariant although this assumption can be relaxed, see section 4.

If A is not competent, i.e. if A is *incompetent*, then he receives only a completely uninformative signal and therefore his posterior assigns probability $1/n$ to all options.

At the end of period t, DM updates his belief that A is competent which is denoted by $\alpha \in [0, 1]$ based on whether the recommendation in period t fitted his needs or not. The initial belief is denoted by α_0. After updating his belief, DM makes a decision whether he wants to *end the game* which means that DM consumes his outside option in period $t + 1$ and there is no further interaction/payoff between DM and A (A will receive a zero payoff in all periods following t in this case while DM get W_O in $t + 1$ and zero ever after). The alternative is to proceed to stage $t + 1$ where the same stage game repeats. A and DM discount future profits. For notational convenience, both are assumed to have the same discount factor $\delta \in (0, 1)$. However, all results continue to hold if A and DM have different discount factors in $(0, 1)$.

DM's outside option is assumed to satisfy

$$1/(n(1 - \delta)) < W_O < p_1/(1 - \delta). \tag{1}$$

That is, DM prefers his outside option to receiving advice by an incompetent adviser but DM prefers to get the best possible advice from a competent adviser to his outside option.

2.1. Strategies, Value Functions and Equilibrium

This paper focuses on Markov strategies where the belief α is the state variable.[8] That is, the players' actions in period t depend only on α and their observations in t. More precisely, A's recommendation depends on α and the identity of the bonus option. DM's decision to continue to period $t + 1$ or to end the game depends on α and on whether the recommendation in t fitted his needs or not. Strategies are assumed

[7] One could imagine an alternative signal technology where A only gets a noisy signal that option 1 is most likely to fit, i.e. $p_1 > p_2 = p_3 = \cdots = p_n > 0$. All results of this paper continue to hold true in this case which is only ruled out to avoid messy case distinctions.

[8] Note that α is commonly known because both DM and A observe whether a recommendation fitted DM's needs and both know the initial belief α_0.

to be measurable functions. A mixed strategy of DM is denoted by (β^+, β^-) with $\beta^+ : [0,1] \to [0,1]$, $\beta^- : [0,1] \to [0,1]$ where $\beta^+(\alpha)$ $(\beta^-(\alpha))$ denotes the probability of continuing if the recommendation fitted (did not fit) DM's needs and his belief is α. It will be convenient to let β^+ and β^- depend on DM's *updated* belief and I will follow this convention. Note, however, that this is equivalent to having DM's strategy depend on the belief at the beginning of the period, i.e. before updating, because DM's strategy depends in any case on whether the recommendation fitted his needs or not. A's strategies will be analyzed below. At this point, a mixed Markov strategy of A can be written as $s : [0,1] \times \{1,\ldots,n\} \to \Delta\{1,\ldots,n\}$ where $s(\alpha, b)$ is a probability distribution over the options which gives for each option the probability that A recommends this option at belief α if the bonus alternative is b. A Markov perfect equilibrium – henceforth *equilibrium* – is a profile of Markov strategies such that none of the two players can profitably deviate in any subgame.

Given such Markov strategies, one can write the players value functions: $V : [0,1] \to \mathbb{R}_+$ denotes A's value function. That is, $V(\alpha)$ is A's expected discounted payoff stream at the start of a period (before the identity of the bonus element is observed) when DM's belief is α. Similarly, $W : [0,1] \to \mathbb{R}_+$ denotes DM's value function.

What conditions are satisfied by equilibrium value functions and equilibrium strategies? The optimal strategy of an incompetent adviser is straightforward: As he cannot influence the probability of satisfying DM's needs (for him all options are equally likely to satisfy DM's needs), it is optimal for him to recommend the bonus option no matter what α is.

Lemma 1. *In equilibrium, the incompetent adviser recommends the bonus option (at any belief α).*

As the strategy of the incompetent adviser, is uninteresting I will in the remainder refer to the competent adviser with "adviser" (A) and with V to this type's value function.

Given some belief α (held at the begin of a period), let us denote by α^+ DM's updated belief – using Bayes' rule – in case the recommendation fits his needs and by α^- his belief if the recommendation does not fit his belief. I will focus on *informative equilibria* with which I mean that $\alpha^+ \geq \alpha^-$. Put differently, a competent adviser in an informative equilibrium gives at least as good advice as an incompetent adviser, i.e. the competent adviser's strategy is at least as informative as suggesting the bonus option.[9] This implies the following technical result which states that in an informative

[9]While non-informative equilibria are economically somewhat non-sensical, they could in principle exist because a competent A has the ability to give worse advice than an incompetent A, e.g. by suggesting option n. If DM expects A to do so (at some belief α), then it might be a best response for A to actually give bad advice (at this belief) in order to improve (!) his reputation: a fitting advice

equilibrium A prefers – abstracting from the bonus payment – a fitting recommendation to a non-fitting recommendation.

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

In the remainder, I concentrate on informative equilibria. Given $\beta^+(\alpha^+)V(\alpha^+) \geq \beta^-(\alpha^-)V(\alpha^-)$, A's equilibrium strategy has to be a *cutoff strategy*: If the bonus option is among options $1, \ldots, k$, A recommends the bonus option and otherwise A recommends option 1. To see this, it is useful to write A's expected utility as

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

where q is the probability that the recommendation satisfies DM's needs (which depends on the specific recommendation) and $\mathbf{1}_{bonus}$ is the indicator function for the bonus option, i.e. it is 1 if A recommends the bonus option and 0 otherwise. Note first that A will recommend either option 1 or the bonus option. Recommending any other option j is dominated by one of these two: If $\beta^+(\alpha^+)V(\alpha^+) > \beta^-(\alpha^-)V(\alpha^-)$, recommending option 1 instead of j increases the probability that the recommendation fits DM's needs (q) and therefore increases the probability of receiving $\beta^+(\alpha^+)V(\alpha^+)$ instead of $\beta^-(\alpha^-)V(\alpha^-)$. If $\beta^+(\alpha^+)V(\alpha^+) = \beta^-(\alpha^-)V(\alpha^-)$, then any recommendation will give the same continuation value but the bonus option gives an additional payoff of 1 today.

Now suppose recommending the bonus option if it is option $k > 1$ yields a (weakly) higher payoff than recommending option 1. Then the same is true if the bonus option is option $j < k$ with $j > 1$. This is true as the probability of fitting DM's needs is higher when the bonus option is j than when it is k. Hence, the expected payoff from recommending the bonus option is higher when it is j than when it is k and therefore it is higher than the expected payoff from recommending option 1. This implies that A uses a cutoff strategy. Note that this argument does not completely rule out mixed strategies: If A is indifferent between two cutoff levels, he could mix over two cutoff strategies. However, the argument above implies that A can only be indifferent between adjacent cutoff levels k and $k+1$.

What is the optimal cutoff chosen by A given an equilibrium value function V? The cutoff level $k \in \{2, \ldots, n-1\}$ is optimal if two conditions are satisfied: First, recommending the bonus option leads to a higher payoff than recommending option 1 if the bonus option is option k. Second, the opposite is true if the bonus option is option

would then be interpreted as being more likely to be given by an incompetent type and therefore reduce α.

$k+1$. The corresponding inequalities are

$$p_k\delta\beta^+(\alpha^+)V(\alpha^+) + (1-p_k)\delta\beta^-(\alpha^-)V(\alpha^-) + 1 \\ \geq p_1\delta\beta^+(\alpha^+)V(\alpha^+) + (1-p_1)\delta\beta^-(\alpha^-)V(\alpha^-)$$

$$p_{k+1}\delta\beta^+(\alpha^+)V(\alpha^+) + (1-p_{k+1})\delta\beta^-(\alpha^-)V(\alpha^-) + 1 \\ \leq p_1\delta\beta^+(\alpha^+)V(\alpha^+) + (1-p_1)\delta\beta^-(\alpha^-)V(\alpha^-).$$

Rearranging these inequalities gives the following result.

Lemma 3. *Let V be an equilibrium value function. The equilibrium cutoff at belief α k_α satisfies*

$$p_1 - p_{k_\alpha} \leq \frac{1}{\delta(\beta^+(\alpha^+)V(\alpha^+) - \beta^-(\alpha^-)V(\alpha^-))} \leq p_1 - p_{k_\alpha+1} \tag{2}$$

if $k_\alpha \in \{2,\dots,n-1\}$. If $k_\alpha = 1$, then

$$\frac{1}{\delta(\beta^+(\alpha^+)V(\alpha^+) - \beta^-(\alpha^-)V(\alpha^-))} \leq p_1 - p_2;$$

and if $k_\alpha = n$ then

$$p_1 - p_n \geq \frac{1}{\delta(\beta^+(\alpha^+)V(\alpha^+) - \beta^-(\alpha^-)V(\alpha^-))}.$$

Note that A is indifferent between cutoffs k and $k+1$ if and only if the second inequality in (2) holds with equality. If A uses the cutoff k, the ex ante probability of satisfying DM's needs is

$$q_k = (n-k)\frac{p_1}{n} + \sum_{j=1}^{k}\frac{p_i}{n}. \tag{3}$$

The updated beliefs α^+ and α^- can then be written as[10]

$$\alpha_k^+ = \frac{\alpha q_k}{(1-\alpha)/n + \alpha q_k} \tag{4}$$

$$\alpha_k^- = \frac{\alpha(1-q_k)}{(1-\alpha)(n-1)/n + \alpha(1-q_k)}. \tag{5}$$

DM's optimal strategy is relatively simple: He ends the game if his expected payoff

[10]If A mixes between cutoffs k with probability q and $k+1$ with probability $1-q$, then $\alpha^+ = q\alpha_k^+ + (1-q)\alpha_{k+1}^+$ etc.

(in $t+1$ and following periods) from continuing the game

$$W(\alpha) = q_{k_\alpha}(1 + \delta W(\alpha^+)) + (1 - q_{k_\alpha})\delta W(\alpha^-) \tag{6}$$

is lower than his outside option W_O. As long as $W(\alpha) > W_O$, it is optimal to continue. DM is willing to use a mixed strategy if and only if $W(\alpha) = W_O$.

3. Results

It is not surprising that DM's equilibrium choice is to end the game for sufficiently low α. By assumption – see (1) – DM's outside option is strictly better than receiving advice from an incompetent adviser forever. If DM is sufficiently convinced to face an incompetent adviser, it will therefore be optimal for him to end the relationship.

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

The more surprising result is that DM will end the game (with some probability) also for sufficiently high α. The intuition for the result is the following: 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 continues even if he gives 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 close to zero dynamic incentives to give good advice. Statically, however, he has an incentive to recommend the bonus alternative as this gives an immediate payoff of 1. Consequently, A will recommend the bonus option no matter what his signal is. But this means that both types of A behave in the same way. This 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. The following theorem states this result more formally and strengthens it for equilibria where both players choose piecewise continuous strategies.

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

If 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 $\varepsilon > 0$ such that DM continues with probability $\beta^-(\alpha) < 1 - \varepsilon$ in case the recommendation does not fit his needs for all $\alpha > \bar{\alpha}$.

The previous theorem established that DM will quit the relationship 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 above should already illustrate that this is not the case (i.e. similar problems as for zero quitting probability emerge also with very small positive quitting probabilities). The following lemma strengthens this intuition by stating that DM quits the 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 T_ε does neither depend on the (initial) belief α nor on the equilibrium.

Lemma 4. *Let $\varepsilon > 0$ and define*[11]

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

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

The previous lemma 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 a relationship is finite and bounded from above by*

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

Theorem 2 is driven by the fact that DM ends the relationship when α is high and not driven by a chance that $\alpha < \underline{\alpha}$ at some point of time. To see this, consider $\alpha \to 1$. For beliefs arbitrarily close to 1, the time until which the belief α could possibly fall below $\underline{\alpha}$ 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.

After deriving properties that hold for all equilibria one might wonder about existence of equilibrium. 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 ends the game for all α and A chooses $k_\alpha = n$ for all α. As in usual cheap talk, informative equilibria can exist as well and the following subsection presents an illustrative example solved numerically. This example will also provide the motivation for a more general result on commitment.

[11] The ceiling $\lceil x \rceil$ is the smallest integer above x and the floor $\lfloor x \rfloor$ is the highest integer below x.

3.1. Example

This section presents an equilibrium of a numerical example. The appendix contains an algorithm to calculate equilibria and an online webappendix contains the specific numerical calculations to solve the here presented example.[12]

There are three options ($n = 3$) and parameter values used in this section are $p_1 = 0.9$, $p_2 = 0.08$, $p_3 = 0.02$, $\delta = 0.95$ and $W_O = 6.75$. The following is an equilibrium:

- If $\alpha < 0.35$, DM ends the game and A gives totally uninformative advice, i.e. $k_\alpha = 3$.
- For $\alpha \in [0.35, 0.75]$, DM continues with probability 1 and A recommends always the most likely option 1, i.e. $k_\alpha = 1$.
- If $\alpha > 0.75$, DM will end the game with a non-zero probability if the last recommendation did not fit his needs. A will mix between the cutoffs $k_\alpha = 2$ and $k_\alpha = 3$. The mixing is such that DM is indifferent between continuing and ending the game, i.e. $W(\alpha) = W_O$. DM's mixing probability is such that A is just indifferent between the cutoff strategies $k_\alpha = 2$ and $k_\alpha = 3$.

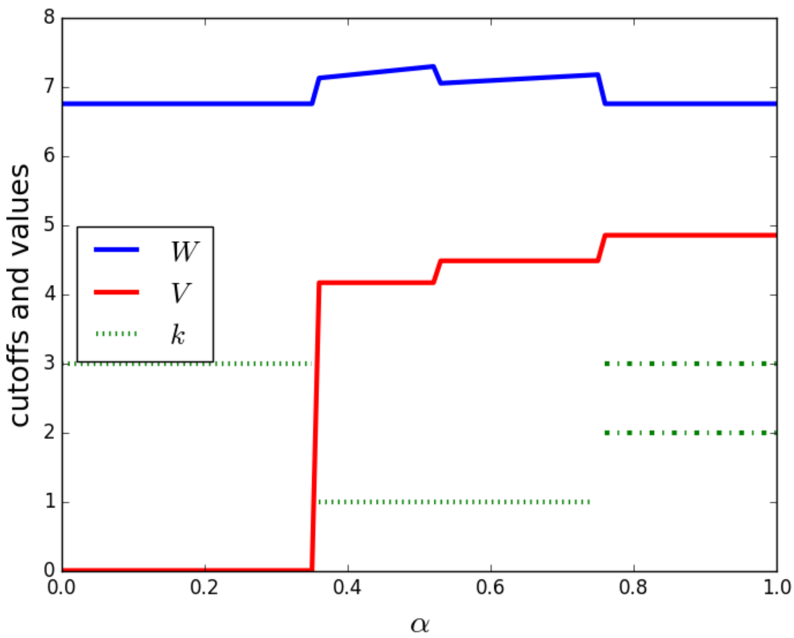

Figure 2: Equilibrium value functions and cutoffs. For $\alpha > 0.75$, A mixes between two cutoffs.

[12]For the code checking the equilibrium conditions and generating the graphs see `https://nbviewer.jupyter.org/urls/schottmueller.github.io/papers/dyn\%20advise\%20reputation/reputation_web1.ipynb`.

The equilibrium strategies and value functions are depicted in figure 2. The example illustrates some general features. DM's value function is equal to his outside option for both very low and very high beliefs α but strictly higher for intermediate beliefs. The mixed strategies for high beliefs are also typical: For high beliefs, DM will have a value equal to his outside option in the next period no matter whether the recommendation fits his needs or not. This follows from $\lim_{\alpha \to 1} \alpha^+ = \lim_{\alpha \to 1} \alpha^- = 1$. But this implies that DM's value today is $\delta W_O + prob$ where $prob$ is the probability of getting a recommendation today that fits his needs. As DM's value today equals his outside option we get that $prob = (1 - \delta)W_O$. This probability is generically not obtained by a pure cutoff strategy and therefore A has to use a mixed strategy for high beliefs. Of course, A uses a mixed strategy only if he is indifferent between the two cutoff levels and this indifference can generically only be ensured if DM uses a mixed strategy.

The example can also illustrate theorem 2 which states that the expected length of the game is bounded by a finite number for all beliefs $\alpha \in (0, 1)$. Figure 3 shows the actual expected length of the game as a function of the belief is less than seven periods for all beliefs.[13] That is, if DM is almost sure that A is of high quality, he expects to ask A for advice only around six times. Figure 4 plots the distribution of game lengths for four different beliefs. It is clear that the probability of asking A for advice more than 20 times is basically zero (conditional on A being actually competent).

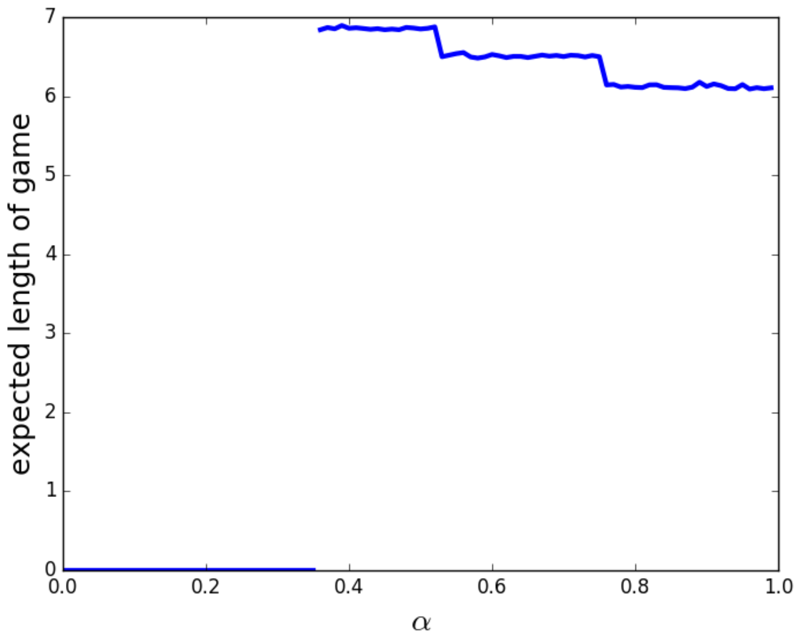

Figure 3: Expected length of game conditional on A being competent (Simulation with 100000 draws per α)

[13]Note that the (expected) length of the games given here already assume that A is competent. The unconditional expected length of the game is shorter.

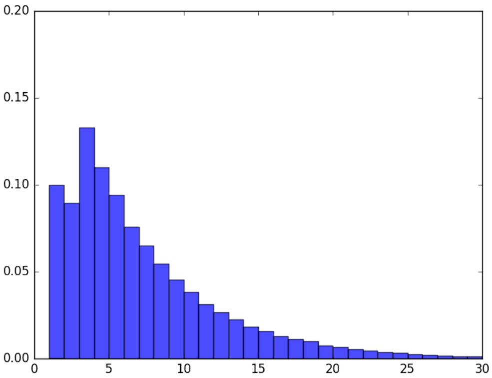

(a) $\alpha = 0.4$

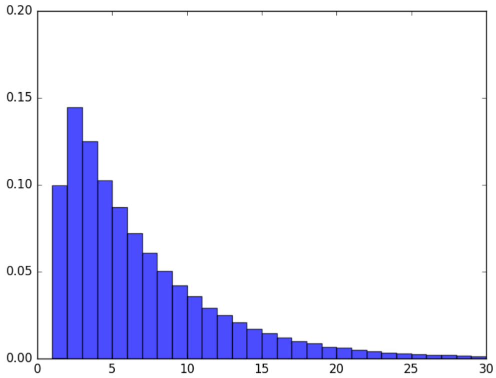

(b) $\alpha = 0.6$

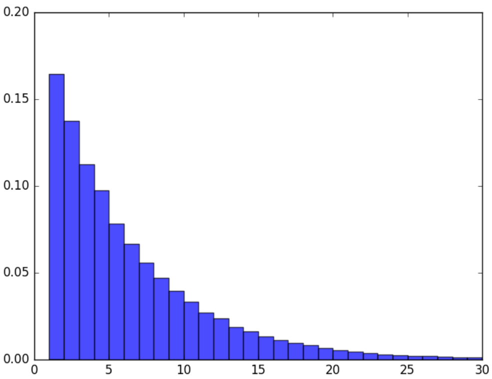

(c) $\alpha = 0.8$

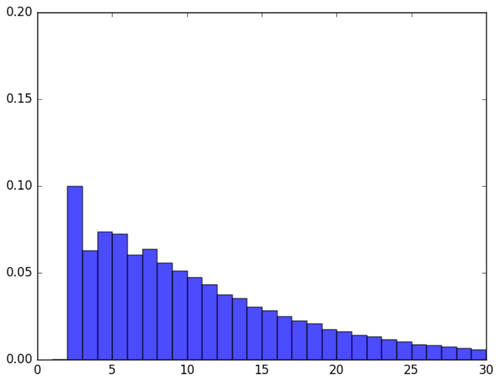

(d) $\alpha = 0.95$

Figure 4: Distribution of game lengths for different beliefs conditional on A being competent (Simulation with 100000 draws per α)

A's value function in figure 2 illustrates a basic commitment problem A faces. Suppose A could commit to the strategy "always recommend option 1". This would give DM the highest possible payoff and imply that DM does not end 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 is 1/3 in each period and therefore the expected payoff of A is $1/(3(1-\delta)) = 20/3$ which is higher than the maximum value of A without commitment. The following proposition states that this situation is not as special as it might appear at first sight. In fact, A's value will be below the commitment value for high levels of α in all equilibria that share the same structure as the example equilibrium, i.e. all equilibria where for high α DM continues for sure if the recommendation fits his needs and ends the game with positive probability if not.

Definition 1. *An equilibrium is called* regular *if (i) players use piecewise continuous strategies,(ii) there exists an $\bar{\alpha} < 1$ such that $\beta^+(\alpha) = 1$ for $\alpha > \bar{\alpha}$ and $\beta^-(\alpha) < 1$ for*

$\alpha > \bar{\alpha}$ and (iii) A uses a mixed strategy for $\alpha > \bar{\alpha}$.

Proposition 2. *In every regular equilibrium,* $\lim_{\alpha \to 1} V(\alpha) < 1/(n(1-\delta))$.

4. Discussion

How can the previous results on dismissal of competent advisers be interpreted? The model allows for two different reasons why DM might want to end the relationship. First, incompetence. If the belief of facing a competent adviser is very low, DM prefers to take his outside option; that is, to fire A. Second, mistrust. This is the case if the belief of facing a competent adviser is high. In this case, A is not afraid of losing his job due to perceived incompetence. But absent this fear A is free to push his own agenda, i.e. he can recommend the bonus option even if it is unlikely to fit DM's needs. Consequently, DM cannot trust A's advice because the fear of losing his job due to incompetence no longer disciplines him. This (fully justified!) lack of trust explains why DM obtains only a value equal to his outside option for high beliefs and why he is willing to fire A in this case.

The previous thoughts also explain why DM obtains his highest values for intermediate beliefs where he is uncertain whether A is competent or not. The uncertainty disciplines A in the sense that A is afraid of being perceived as an ignorant type and being consequently fired due to incompetence. To avoid this scenario, he is willing to give good advise, i.e. he is willing to forgo bonuses today in order to keep the relationship which promises future bonuses. It might – at first glance – appear that this hope for future bonuses is in vain as A can also be fired at high beliefs. This view is, however, misguided for several reasons: First, there is a chance that tomorrow the bonus action coincides with an option that is likely to fit DM's needs, i.e. the tension between wanting to prolong the relationship and getting a bonus could – by chance – be smaller in the next period(s). Second, a fitting recommendation today might make it possible not to be fired even after a non-fitting recommendation in the period afterwards. But then a fitting recommendation today would allow A to collect – at least – two bonuses (in the two consecutively following periods) while a non-fitting recommendation today might lead to an immediate end of the relationship. Third and most importantly, even if a fitting recommendation leads into the region where $\beta^{-}(\alpha) < 1$, A's value at these beliefs might nevertheless be quite high: DM only ends the relationship at these beliefs with some probability and even that only when receiving a non-fitting recommendation; this risk might be small and can easily be outweighed by the decreased likelihood of being dismissed for incompetence anytime soon.

Another way to interpret the results is in terms of hidden information and moral hazard. Hidden information is part of the model as the adviser's type is private. It explains why advisers are fired for low beliefs. For high beliefs, however, the reason for the firing is moral hazard: While the adviser would have the ability to be useful, he decides not be and pursues his private benefit instead.

The results of this paper establish an inefficiency: The game is expected to end in finite time although the advice relationship lasts forever in the first best scenario. The inefficiency arises due to the assumption that A cannot commit to a strategy because his signal is private. That is, an announcement like "I will always recommend the option most likely to fit DM's needs." is not credible because DM cannot check whether A actually sticks to his announcement. Of course, the private nature of A's signal captures exactly the reason why DM needs to get advice and is therefore indispensable. Efficiency gains could be realized if A was able to truthfully reveal his signal. However, this seems unrealistic in many applications. In particular, A will often be able to manipulate the revealed signal which renders the revelation useless.

It is perhaps not surprising that, say, a consumer buys the wrong investment products if his financial adviser receives a bonus for selling certain products. The inefficiency established in this paper is, however, somewhat more subtle. Not only does the consumer buy the wrong products, he will also switch to worse financial advisers after some time or not consume any advice – depending how the outside option in the current model is interpreted. This is an additional inefficiency that arises inevitably if one seriously considers the dynamic nature of the problem. Dynamics should, of course, be considered given that no meaningful advice is possible in the static model as pointed out in the introduction.

One possible solution to the problem would be to resolve the underlying differences 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 compensation of analysts cannot depend on investment banking activities.[14] However, as many commentators pointed out, analysts still receive trading commissions meaning that the problem of non-aligned incentives between analysts and customers was only mitigated but not resolved. Trading commissions are, of course, an instrument to resolve 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. In case of managers and their (subordinate) advisers, the bonus could be non-monetary (e.g. concerns for the own career) which makes it again hard to eliminate. Nevertheless,

[14]See `https://www.sec.gov/news/speech/factsheet.htm` for more information.

the inefficiency established in this paper can be a rational for regulations as the one mentioned above and also for the existence of independent, non-profit consumer organizations. The underlying conflict in preferences does not occur in the case of these organizations and meaningful advice becomes feasible.

The paper explains why most advice relationships are short lived. This 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 therefore he might be willing to give better advice today in order not to risk future bonus payments. Of course, the opposite also holds: If $\delta \to 0$, the bounds get very tight. In fact, T' and $\bar{T}$ are 0 for $\delta < \underline{\delta}$ where the threshold $\underline{\delta}$ is strictly positive. That is, DM will end the game without getting advice once if there is heavy discounting. The reason is that the game with heavy discounting is similar to the static game where A will always recommend his bonus option and meaningful advice is therefore impossible.

Finally, 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 consider non-Markovian strategies and equilibria. In this case, theorem 1 can no longer be stated as DM's strategy will depend not only on α but possibly on the complete history of the game. It is therefore not clear how a result like "advisers who are believed to be competent are fired" can be stated if strategies do not directly depend on the belief of being competent. In contrast to this, the results on the (expected) length of the game, i.e. lemma 4 and theorem 2, will continue to hold in every – also non-Markovian – equilibrium. The reason is that the proof uses only a general non-deviation constraint that has to be satisfied in every equilibrium. Third, one might wonder about a signal technology that is not constant over time. That is, the posterior $p_1, \ldots, p_n$ might depend on the period. Again theorem 1 is then impossible to state as strategies will then naturally not only depend on the belief but also on the time period. The results on the (expected) length of the game, however, still hold true if one substitutes p_n by $\inf(\{p_n^t\})$ where p_n^t is p_n in period t and the infimum is assumed to be strictly larger than 0. With this adjustment the proofs of lemma 4 and 2 will go through and the results hold. If the posterior $p_1, \ldots, p_n$ depended not on the time period t but on the belief α, then even theorem 1 could still be stated and proven as long as $\inf(\{p_n(\alpha)\}) > 0$.

One might also wonder whether it is possible to restrict the strategy space such that DM's decision is allowed to depend on his belief α only (and not on whether the recommendation in the current period fitted his needs or not); that is, $\beta^+(\alpha) = \beta^-(\alpha)$. Unfortunately, no informative equilibrium with piecewise continuous strategies apart from the babbling equilibrium exists in this restricted class. To see this, consider $\beta(\alpha)\,(V(\alpha^+) - V(\alpha^-))$ as $\alpha \to 1$. With piecewise continuous strategies V will be continuous for all α above a certain threshold and from Bayes' rule, see (4) and (5), it follows that $(V(\alpha^+) - V(\alpha^-))$, and therefore also $\beta(\alpha)\,(V(\alpha^+) - V(\alpha^-))$, converge to 0 as $\alpha \to 1$. Consequently, $k(\alpha) = n$ for sufficiently high α by lemma 3 which implies that DM has to end the game for all α above a certain threshold $\hat{\alpha} < 1$. Consequently, $V(\alpha) = 0$ for α sufficiently high. But this implies that the equilibrium cannot be informative unless $V(\alpha) = 0$ for all $\alpha \in [0, 1]$: Otheriwse, there would have to be a belief α where $V(\alpha^-) > 0$ but $V(\alpha^+) = 0$ and a competent adviser would then have an incentive to give worse advice than an incompetent type. Of course, $V(\alpha) = 0$ for all α implies that DM never asks A for advice, i.e. this value function belongs to the babbling equilibrium. Hence, the intuitive assumption that DM's equilibrium strategy depends on whether he just got fitting advice or not is essential to construct equilibria with some information transmission.

5. Conclusion

This paper analyzes the questions why advisers are fired. Two reasons 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 implies that the decision maker benefits most 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 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 limited although advice by a qualified adviser is efficient. The presence of private benefits for the adviser, like bonus payments, does therefore not only lead to bad advice, it also implies that people 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.[15] While this literature establishes empirically that more reputable analysts give on average better predictions, this result is based on assuming either a binary or a linear functional form for this relationship. The model of this paper, however, suggests a non-monotonic (possibly inversely U-shaped) relationship.[16] On the theoretical side, exploring alternative setups provides opportunities for further research. For example, competition between advisers is only part of the model if one interprets the outside option as the value that can be received from advice by another adviser. If simultaneous competition between several advisers is modeled explicitly, switching forth and back between advisers might become a viable strategy. This and other possibilities are beyond the scope of the current paper and left for future research.

[15] "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.

[16] 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 shows that this argument, though plausible at first sight, might not be true in equilibrium.

Appendix

5.1. Calculating (simple) equilibria for $n = 2$

This section shows how to calculate regular equilibria. To simplify matters, I will (a) concentrate for most of the section on the case where there are only two options, i.e. $n = 2$, and (b) focus on equilibria in which neither player mixes at $\alpha < \bar{\alpha}$ where $\bar{\alpha}$ is as in definition 1. This implies that the equilibria of this section have the structure described in figure 5 (where the statement concerning V' will be derived below). I will derive an equilibrium candidate that has this structure (namely value functions V and W) and then to formulate conditions under which this equilibrium candidate is indeed an equilibrium. This will be done for fixed $\underline{\alpha}$ and $\bar{\alpha}$. That is, one should think of repeating this derivation and check for different values of $\underline{\alpha}$ and $\bar{\alpha}$ whether the derivation leads to an equilibrium.

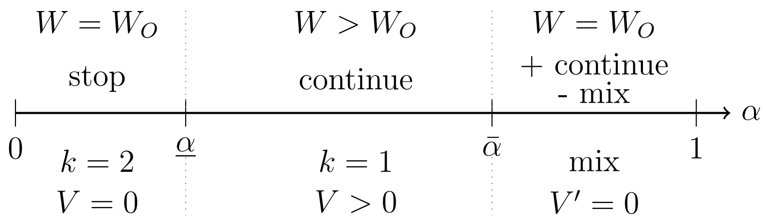

Figure 5: Equilibrium structure: First "row" indicates DM's value function, second row DM's strategy, third row (below the axis) A's strategy and fourth row A's value function.

First consider $\alpha > \bar{\alpha}$. As A mixes, he has to be indifferent between $k = 1$ and $k = 2$. Hence,

$$\beta^-(\alpha) = \frac{V(\alpha^+) - \frac{1}{\delta(p_1 - p_2)}}{V(\alpha^-)}. \tag{7}$$

As A is indifferent between his equilibrium strategy and using the cutoff 2 with probability 1, his value can be written as $V(\alpha) = q_2 V(\alpha^+) + (1 - q_2)\beta^-(\alpha)V(\alpha^-) + 1$. Plugging $\beta^-(\alpha)$ as stated above into A's value function gives

$$V(\alpha) = \delta V(\alpha^+) - \frac{1 - q_2}{p_1 - p_2} + 1 = \delta V(\alpha^+) - \frac{1/2}{2p_1 - 1} + 1$$

where the last equality utilizes that $n = 2$ implies $q_1 = p_1$, $q_2 = 1/2$ and $p_2 = 1 - p_1$. Note that the previous equation is solved by a constant V where (of course only for $\alpha > \bar{\alpha}$)

$$V(\alpha) = V^* \equiv \frac{1}{1 - \delta} - \frac{1}{2(1 - \delta)(2p_1 - 1)}.$$

Second, consider $\alpha < \underline{\alpha}$. Note that $k = 2$ implies that $\alpha^+ = \alpha^- = \alpha$, i.e. there is no updating because both types of A use the same strategy of always recommending their bonus option. Given the other player's strategy and the fact that there is no updating, both A and DM clearly play a best response.

Third, consider $\alpha \in [\underline{\alpha}, \bar{\alpha}]$. Given A's strategy of using cutoff $k = 1$, the updated beliefs are

$$\alpha^+ = \frac{\alpha p_1}{(1-\alpha)/2 + \alpha p_1} = \frac{1}{1 + \frac{1-\alpha}{\alpha}\frac{1}{2p_1}} \tag{8}$$

$$\alpha^- = \frac{\alpha(1-p_1)}{(1-\alpha)/2 + \alpha(1-p_1)} = \frac{1}{1 + \frac{1-\alpha}{\alpha}\frac{1}{2(1-p_1)}}. \tag{9}$$

A's value function for $\alpha \in (\underline{\alpha}, \bar{\alpha})$ has to satisfy

$$V(\alpha) = \frac{1}{2} + p_1\delta V(\alpha^+) + (1-p_1)\delta V(\alpha^-) \tag{10}$$

where α^+ and α^- are as stated above. The main difficulty is to find a function V that satisfies this equation. Define the mapping $\tilde{V}$ that assigns to each function $V : [\underline{\alpha}, \bar{\alpha}] \to \mathbb{R}_+$ the function $\tilde{V}$ defined by

$$\tilde{V}(V)(\alpha) = \frac{1}{2} + p_1\delta V(\alpha^+) + (1-p_1)\delta V(\alpha^-)$$

where α^+ and α^- are as stated in (8) and (9) (using $V(\alpha^-) = 0$ for $\alpha^- < \underline{\alpha}$ and $V(\alpha^+) = V^*$ for $\alpha^+ > \bar{\alpha}$). It is straightforward to show that $\tilde{V}$ is a contraction.[17] Applying the contraction mapping theorem yields that $\tilde{V}$ has a unique fixed point which implies that there is a unique increasing V that satisfies (10). The contraction mapping theorem also gives a way to compute this V: Starting from any bounded function, applying $\tilde{V}$ repeatedly will lead to a sequence of value functions converging to V.

The steps above gave us a solution candidate for V (given $\underline{\alpha}$ and $\bar{\alpha}$). Before checking whether this candidate is part of an equilibrium, we derive an equilibrium candidate for W. By the structure of the equilibrium, $W(\alpha) = W_O$ for $\alpha < \underline{\alpha}$ and $\alpha > \bar{\alpha}$. For $\alpha \in [\underline{\alpha}, \bar{\alpha}]$, W has to satisfy

$$W(\alpha) = (\alpha p_1 + (1-\alpha)/2)(1 + \delta W(\alpha^+)) + (1 - \alpha p_1 - (1-\alpha)/2)\delta W(\alpha^-)$$

where α^+ and α^- are as in (8) and (9). The mapping $\tilde{W}$ that assigns to every bounded function W on $[\underline{\alpha}, \bar{\alpha}]$ the function $\tilde{W}(W)$ given by

$$\tilde{W}(W)(\alpha) = (\alpha p_1 + (1-\alpha)/2)(1 + \delta W(\alpha^+)) + (1 - \alpha p_1 - (1-\alpha)/2)\delta W(\alpha^-)$$

[17] For this, we equip the space of bounded functions on $[\underline{\alpha}, \bar{\alpha}]$ with the sup norm.

is again a contraction. By the contraction mapping theorem, a unique fixed point of $\tilde{W}$ exists and this fied point can be obtained by iterating $\tilde{W}$. This fixed point is the candidate solution for W.

When is this candidate actually an equilibrium? The steps above ensure that the candidate V is consistent with DM's strategy and also that A is indifferent between $k = 1$ and $k = 2$ for $\alpha > \bar{\alpha}$. It is, however, unclear whether A's strategy is optimal for $\alpha \in (\underline{\alpha}, \bar{\alpha})$. Following lemma 3, the candidate V is part of an equilibrium only if

$$0 \leq \frac{1}{\delta(V(\alpha^+) - V(\alpha^-))} \leq 2p_1 - 1 \quad \text{for } \alpha \in [\underline{\alpha}, \bar{\alpha}].$$

Furthermore, V has to be non-negative everywhere. While this is obviously true for $\alpha < \bar{\alpha}$, it is unclear for V^* (note that we did not track the restriction $\beta^-(\alpha^-) \in [0, 1]$ when deriving V^*). In fact, $V^* \geq 0$ if and only if

$$p_1 \geq 3/4.$$

This means that the type of equilibria discussed here exists only if the information technology is sufficiently informative as one would have expected. For W, it is necessary to check that $W \geq W_O$. By construction, this is true for $\alpha < \underline{\alpha}$ and $\alpha > \bar{\alpha}$ but needs to be checked in for $\alpha \in [\underline{\alpha}, \bar{\alpha}]$.

Finally, it is necessary to check whether the mixing is feasible for all $\alpha > \bar{\alpha}$. That is, is $\beta^-(\alpha) \in [0, 1]$ as defined in (7) and is there a $s(\alpha) \in [0, 1]$ (which is interpreted as the probability that A uses $k = 1$) for every $\alpha > \bar{\alpha}$ such that DM is indifferent between his outside option and continuing? Note that if $s(\alpha) = 0$, the quality of the advice DM is getting is the same as getting advice from an incompetent adviser forever. By assumption, the outside option is better than this. Hence, it is only necessary to check that for $s(\alpha) = 1$ continuing is optimal for DM: By continuity, there will then be a $s(\alpha) \in (0, 1]$ where DM is indifferent. Note furthermore that continuing is in case $s(\alpha) = 1$ is always optimal for α sufficiently high: By assumption getting best advice ($k = 1$) from a competent adviser is better than the outside option. As DM's value in the next period is at least the outside option, continuing is therefore optimal if $s(\alpha) = 1$ and α close to 1. Hence, existence of a suitable $s(\alpha)$ will be unproblematic unless $\bar{\alpha}$ is too low.

To summarize, it is computationally simple to obtain a candidate equilibrium (calculating the fixed points of two contractions) for given $\underline{\alpha}$ and $\bar{\alpha}$. A candidate equilibrium is an equilibrium if

- $p_1 \geq 3/4$,
- $0 \leq \frac{1}{\delta(V(\alpha^+) - V(\alpha^-))} \leq 2p_1 - 1 \quad \text{for } \alpha \in [\underline{\alpha}, \bar{\alpha}]$,

- $W(\alpha) \geq W_O$ for $\alpha \in [\underline{\alpha}, \bar{\alpha}]$,

- $0 \leq \beta^-(\alpha) \leq 1$ for $\alpha > \bar{\alpha}$ where $\beta^-(\alpha)$ is defined as in (7),

- continuing is optimal if $s(\theta) = 1$: $\alpha \left[p_1(1 + \delta W(\alpha^+)) + (1 - p_1)\delta W(\alpha^-)\right] + (1 - \alpha)\left[(1 + \delta W(\alpha^+))/2 + \delta W(\alpha^-)/2\right] \geq W_O$ for all $\alpha > \bar{\alpha}$ (where $W(\alpha^+) = W_O$ by the structure of the equilibrium).

Again these conditions are easy to check numerically.

For $n > 2$, in principle the same procedure can be used to produce candidate equilibria. However, some additional difficulties arise. In particular, one might be interested in equilibria that have cutoff $k = 2$ for some $\alpha \in [\underline{\alpha}, \bar{\alpha}]$ and cutoff $k = 1$ for others. This makes it impractical to fix the strategies of the players before deriving the candidate value functions. One could extend the procedure for finding a candidate V in the following way: Starting from some arbitrary V^0 determined A's optimal cutoffs given V^0 (using lemma 3). Then use these cutoffs and V^0 in the $\tilde{V}$ operator to get V^1 and iterate this procedure. The problem with this procedure is that as soon as A's strategy is endogenous, i.e. new cutoffs are computed in each iteration, $\tilde{V}$ is no longer a contraction. This has two implications: First, it might have several fixed points or no fixed points at all. Second, even if it has a fixed point the iteration is not guaranteed to converge. If, however, it converges to a fixed point a candidate equilibrium emerges and the remaining feasibility conditions can be checked for this candidate equilibrium as above.

5.2. Proofs

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. 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. $\square$

Proof of proposition 1: DM's expected payoff from continuing is bounded from above by $\alpha p_1/(1-\delta) + (1-\alpha)/(n(1-\delta))$. For α sufficiently low (but strictly higher than 0), this upper bound is less than W_O as $W_O > 1/(n(1-\delta))$ by (1). $\square$

Proof of theorem 1: The first part is proven by contradiction and follows the argument in the main text. Suppose the statement was not true; i.e. suppose that there was an equilibrium such that for no sequence $(\alpha_i)_{i=1}^{\infty}$ converging to 1 there exists $\varepsilon > 0$ such that DM ends the game with at least ε probability at each element of the sequence.[18] This implies that in this equilibrium for every $\varepsilon' > 0$ there exists an $\bar{\alpha}_{\varepsilon'} < 1$

[18] It is immaterial for this proof whether DM ends the game only when receiving a non-fitting

such that for all $\alpha \geq \bar{\alpha}_{\varepsilon'}$ DM continues with probability greater than $1-\varepsilon'$. Note that $\alpha - \alpha^-$ converges to zero as $\alpha \to 1$ (for any strategy A employs). 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 optimal to choose the pure strategy $k_{\alpha_{T\varepsilon'}} = n$: The reason is that A at belief $\alpha_{T\varepsilon'}$ when observing that the bonus option is option n 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 n would lead to a payoff of at most $\delta/(1-\delta) = \delta + \delta^2 + \delta^3 + \ldots$. 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. This establishes the claim that A chooses the pure cutoff strategy $k_{\alpha_{T\varepsilon'}} = n$ for T high and $\varepsilon > 0$ small enough. But then $\alpha_{T\varepsilon'} = \alpha^+_{T\varepsilon'} = \alpha^-_{T\varepsilon'}$ and therefore $W(\alpha_{T\varepsilon'}) = 1/(n - n\delta) < W_O$; 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 later). Note that whenever A plays the pure strategy $k_\alpha = n$ 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 $1/(n - n\delta) < W_O$.

Choose $k \in \{1, 2, \ldots, n-1\}$ such that for every $\bar{\alpha} < 1$ we have $k = k_\alpha$ (or more generally for the case of mixed strategies: $k_\alpha = k$ with positive probability) for some $\alpha > \bar{\alpha}$. If no such k exists, then $k_\alpha = n$ for sufficiently large α and the previous paragraph implies that the last claim of the theorem is true. I will now show that this claim is also true if such a $k \in \{1, 2, \ldots, n-1\}$ exist. Let A_ε be the set of α such that $k_\alpha = k$ (with positive probability) and $\beta^-(\alpha) > 1 - \varepsilon$. If the last claim of the theorem holds, then $A_\varepsilon \cap (\bar{\alpha}, 1)$ is empty for $\varepsilon > 0$ small enough and $\bar{\alpha} < 1$ large enough. Suppose this is not the case, i.e. suppose $A_\varepsilon \cap (\bar{\alpha}, 1)$ is non-empty for all $\varepsilon > 0$ and $\bar{\alpha} < 1$. For $\alpha \in A_\varepsilon$, the difference $\beta^+(\alpha)V(\alpha^+) - \beta^-(\alpha)V(\alpha^-)$ is bounded from below by lemma 3 and $k < n$. This and $\beta^-(\alpha) > 1 - \varepsilon$ implies that $V(\alpha^+) - V(\alpha^-)$ is bounded from below as well for $\varepsilon > 0$ sufficiently small and $\alpha \in A_\varepsilon$. Furthermore, if we concentrate on $\alpha \in A_\varepsilon$, α^+ and α^- are strictly increasing in α as $k_\alpha = k$ for all $\alpha \in A_\varepsilon$. 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_ε 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

recommednation or not. The reason is that A's signal is noisy. Hence, even when recommending option 1 there is a probability $1 - q_1 > 0$ of not fitting. 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 $\varepsilon = (1 - q_1)\tilde{\varepsilon}$.

$\alpha^- > a_i^+$ for all $\alpha > \hat{a}_i^+$ (this follows as $\alpha - \alpha^-$ converges to zero as α converges to 1 and the assumption that $A_\varepsilon \cap (\bar{\alpha}, 1)$ 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 then implies that V is a function of unbounded total variation. This, however, contradicts the piecewise continuity of V. Hence, $A_\varepsilon \cap (\bar{\alpha}, 1)$ has to be empty for $\bar{\alpha} < 1$ high enough and $\varepsilon > 0$ 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: If $k_\alpha = n$, then $\alpha^+ = \alpha^- = \alpha$ and $W(\alpha) < W_O$, i.e. DM ends the relationship immediately and any $T_\varepsilon > 0$ will give the result.

Take α such that $k_\alpha < n$ with positive probability. Then the probability that DM ends the game within T' periods (when DM sticks to his equilibrium strategy while A might not) has to be at least ε' for some $\varepsilon' > 0$ (taking a high T' and a sufficiently small $\varepsilon' > 0$). That is, there is at least one path of (up to) T' hits and misses such that the probability that DM ends the game along this path is at least $\varepsilon' > 0$. To see this, note that $k_\alpha < n$ implies that A (weakly) prefers to recommend option 1 to recommending the bonus option if the bonus option is option n. This is clearly not the case if no ε' and T' as described above exist.

Next, I will show that we can use $\varepsilon' = (1-\delta)/2$ and $T' = \lfloor 2\log(1-\delta)/\log(\delta) - 1 \rfloor$. Suppose otherwise, i.e. suppose there is no path of hits and misses 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 has an incentive to deviate when the bonus option is option n. 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:

$$\begin{aligned}
T' = \left\lfloor 2\frac{\log(1-\delta)}{\log(\delta)} - 1 \right\rfloor &\leq 2\frac{\log(1-\delta)}{\log(\delta)} - 1 \\
\Leftrightarrow (T'+1)\log(\delta) &\leq 2\log(1-\delta) \\
\Leftrightarrow \delta^{T'+1} &\leq (1-\delta)^2 \\
\Leftrightarrow \frac{\delta^{T'+1}}{1-\delta} &\leq 1-\delta \\
\Leftrightarrow \delta\left(1+\delta^{T'+1}+\delta^{T'+2}+\dots\right) &\leq 1 \\
\Leftrightarrow (1-\delta)\left(\delta+\delta^2+\dots+\delta^{T'}\right)+\delta^{T'+1}+\delta^{T'+2}+\dots &\leq 1 \\
\Leftrightarrow \delta+\delta^2+\dots &\leq 1+\delta\left(\delta+\delta^2+\dots+\delta^{T'}\right) \\
\Leftrightarrow \frac{\delta}{1-\delta} &\leq 1+(1-2\varepsilon')\left(\delta+\delta^2+\dots+\delta^{T'}\right)
\end{aligned}$$

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 of the supposed equilibrium payoff. This establishes that ε' and T' have the desired property.

As the path of hits and misses 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. $p_1 < 1$ and $p_n > 0$), 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 under equilibrium play (which can be chosen independent of the specifics of the equilibrium and the belief depending only on the probabilities $p_1, \dots, p_n$, e.g. $\gamma = p_n^{T'}$ 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 $\gamma = p_n^{T'}$ and T', ε' as derived above, for example, yields

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

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

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

equals ε (which again will increase the expectation as $\tilde{T}_\varepsilon \geq T_\varepsilon$) for $\tilde{T}_\varepsilon > T'$. 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 $B = T'/\log\left(1 - p_n^{T'}\varepsilon'\right)$. 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

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

□

Proof of proposition 2: Without loss of generality, choose $\bar{\alpha}$ high enough such that strategies and value functions are continuous for $\alpha > \bar{\alpha}$. As DM uses a mixed strategy, his value $W(\alpha)$ has to equal W_O for $\alpha > \bar{\alpha}$. Therefore, A mixes over two cutoff levels $\tilde{k}$ and $\tilde{k}+1$ such that using cutoff $\tilde{k}$ for sure would result in a DM value above W_O while using $\tilde{k}+1$ for sure would result in a DM value below W_O. A is indifferent between the cutoffs $\tilde{k}$ and $\tilde{k}+1$ only if

$$\beta^-(\alpha) = \frac{V(\alpha^+) - \frac{1}{\delta(p_1 - p_{\tilde{k}+1})}}{V(\alpha^-)}.$$

As A is indifferent between his equilibrium strategy and using the cutoff $\tilde{k}+1$ with probability 1, his value can be written as $V(\alpha) = q_{\tilde{k}+1}\delta\beta^+(\alpha)V(\alpha^+) + (1-q_{\tilde{k}+1})\delta\beta^-(\alpha)V(\alpha^-) + (\tilde{k}+1)/n$. Plugging $\beta^-(\alpha)$ and $\beta^+(\alpha) = 1$ into A's value function gives

$$V(\alpha) = \delta V(\alpha^+) - \frac{1 - q_{\tilde{k}+1}}{p_1 - p_{\tilde{k}+1}} + \frac{\tilde{k}+1}{n}.$$

Since V is continuous and $\alpha^+ \to \alpha$ as $\alpha \to 1$, the limit result $\lim_{\alpha\to1} V(\alpha) = \lim_{\alpha\to1} V(\alpha^+)$ has to hold. Consequently, the previous equation can be solved in the limit for $V(\alpha)$:

$$\lim_{\alpha\to1} V(\alpha) = \frac{1}{1-\delta}\left(-\frac{1 - q_{\tilde{k}+1}}{p_1 - p_{\tilde{k}+1}} + \frac{\tilde{k}+1}{n}\right).$$

It is now shown that $\lim_{\alpha \to 1} V(\alpha) < 1/n$, i.e. that

$$\frac{\tilde{k}}{n} < \frac{1 - q_{\tilde{k}+1}}{p_1 - p_{\tilde{k}+1}}.$$

Plugging in the definition of $q_{\tilde{k}+1}$ in (3) and multiplying through by $(p_1 - p_{\tilde{k}+1})n$ gives

$$\begin{aligned}
\tilde{k}(p_1 - p_{\tilde{k}+1}) &< n - (n - \tilde{k} - 1)p_1 - \sum_{i=1}^{\tilde{k}+1} p_i \\
\Leftrightarrow 0 &< n(1 - p_1) - \sum_{i=2}^{\tilde{k}} (p_i - p_{\tilde{k}+1}) \\
\Leftrightarrow 0 &< (n - \tilde{k} + 1)(1 - p_1) + \sum_{i=2}^{\tilde{k}} (1 - p_1 - p_i + p_{\tilde{k}+1})
\end{aligned}$$

which is obviously true as $1 - p_1 - p_i \geq 0$ for $i = 2, \ldots$. $\square$

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.

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.

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.

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.

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