Can a Bonus Overcome Moral Hazard? An Experiment on Voluntary Payments, Competition, and Reputation in Markets for Expert Services

Vera Angelova (TU Berlin)
Tobias Regner (Friedrich Schiller University Jena)

Discussion Paper No. 26

March 29, 2017

Can a Bonus Overcome Moral Hazard?
An Experiment on Voluntary Payments, Competition, and Reputation in Markets for Expert Services[‡]

Vera Angelova[*] and Tobias Regner[**]

Abstract

Interactions between players with private information and opposed interests are often prone to bad advice and inefficient outcomes, e.g. markets for financial or health care services. In a deception game we investigate experimentally which factors could improve advice quality. Besides advisor competition and identifiability we add the possibility for clients to make a voluntary payment, a bonus, after observing advice quality. We observe a positive effect on the rate of truthful advice when the bonus creates multiple opportunities to reciprocate, that is, when the bonus is combined with identifiability (leading to several client-advisor interactions over the course of the game) or competition (allowing one advisor to have several clients who may reciprocate within one period).

JEL classification: C91, D03, D82, G20, I11

Keywords: asymmetric information; principal–agent; expert services; deception game; sender–receiver game; reciprocity; reputation; experiments; voluntary payment; competition

[*]Technische Universität Berlin, Berlin, Germany, e-mail: vera.angelova@tu-berlin.de.

[**]Friedrich Schiller Universität, Jena, Germany, e-mail: tobias.regner@uni-jena.de.

1 Introduction

Moral hazard on the financial market is detrimental for consumers. Empirical evidence from the US shows that mutual funds offering higher broker commissions attract the most investments. However, higher commissions are related to lower investment performance (Christoffersen et al., 2013). Clients in Germany lose 50 billion euros per year due to misleading financial advice ('Die Welt', 2012). Mullainathan et al. (2012) report that retail financial advisors tend to give self-serving portfolio recommendations. An audit study focusing on the Indian life insurance market reports that life insurance agents recommend strictly dominated products which yield high commissions in up to 90% of the cases (Anagol et al., 2013). These market inefficiencies are due to asymmetric information (uninformed clients) and commission steering by funds (see Inderst and Ottaviani, 2012). In similar fashion, the health care sector which accounts for 15% of GDP in OECD countries (OECD, 2016) suffers from moral hazard and efficiency losses caused by information asymmetries.

Which factors could contribute to experts providing better advice and, in turn, to increased market efficiency?[1] We design a laboratory experiment to analyze the stylized relationship between expert advisors and clients in a controlled setting. As our experimental framework we use a deception game (Gneezy, 2005). We augment it with market forces (competition, the possibility to build reputation) and allow a voluntary action of the client, for instance, a bonus payment (at the end of the transaction after feedback about quality of advice has been provided).

The key innovation of our design is to test whether a voluntary component can be a remedy against moral hazard, on its own and in interaction with instruments that have been studied before (competition and reputation). Huck et al. (2012) use a binary-choice trust game to analyze experience goods markets. They conclude that reputation based on quality provided

---

[1]See Angelova and Regner (2013) for the connection between advice quality and efficiency in the market of financial intermediaries.

in the past enhances trust and that competition reinforces this effect. Dulleck et al. (2011) analyze the richer framework of credence goods.[2] They find little effect of reputation and no effect of competition on undertreatment.

We use a 2x2x2 (competition, identifiability, option to pay a bonus) between-subjects design and model advice as an experience good. Its quality is unknown ex-ante but ex-post the client finds out whether the advice was good or bad. We focus on markets where clients' access to advisors' past behavior is limited to own past observations, in contrast to centralized market platforms that allow easy access to an online history of transactions (e.g., ebay or Amazon). Moreover, our study addresses types of advice where taking the outside option is not plausible: for instance, a medical treatment that needs to be taken or advice for a required investment. This distinguishes our setting from online purchases where a credible outside option commonly exists (buying at the local shop, albeit at a higher price).

Without competition, we find a significant increase in the rate of truthful advice when a bonus can be given and advisors are identifiable. With competition, the rate of truthful advice is higher when a bonus can be given or when advisors are identifiable. Our results are in line with Huck et al. (2012) as the combination of competition and reputation concerns reduces opportunistic behavior in a setting where cheating is detectable. Results from our bonus treatments indicate that a voluntary component can also lead to a reduction of advisors' cheating. Multiple opportunities to reciprocate are a necessary condition. This can be achieved in the time or client dimension. Identifiability leads to several client-advisor interactions over the course of the game and competition allows one advisor to have several clients who may reciprocate within one period. In reality, the bonus could be thought of as any voluntary act that is costly to the client but would benefit the advisor as, e.g., spreading the word about the advisor on an online feedback platform.

---

[2]Experience goods can be seen as a subset of credence goods. Their common feature is the possibility of undertreatment, i.e. providing a lower quality than required. Besides, credence goods also allow for overtreatment and overcharging.

The next section discusses the related literature. In section 3 we explain our experimental set-up, relate it to existing studies, and state our behavioral predictions. In section 4 we present the results and discuss them. Section 5 concludes.

2   Related literature

We begin with a review of experience/credence goods studies, focusing on the two experiments that are closest to ours, and then proceed with the literature on deception that is relevant in our context.

Huck et al. (2012), henceforth HLT, use a repeated binary-choice trust game to analyze the effects of reputation and competition in a market for an experience good. They vary the extent with which trustors are informed about past behavior of trustees. There is either no, private (i.e. only about trustees a trustor has interacted with in the past) or public information (i.e. about all past interactions of all trustees). Moreover, trustors are either exogenously matched with a trustee (no-competition-treatment) or they can choose their preferred trustee based on her reputation (competition-treatment). HLT find that reputation enhances trust (but no difference between private and public information) and that reputation combined with competition eliminates the trust problem almost completely.[3]

Dulleck et al. (2011), henceforth DKS, study the effect of institutions (liability, verifiability), market forces (competition, reputation), and combinations of these on the provision of credence goods. In DKS's setting clients are uncertain about the quality they need. Sellers know what clients need but can offer either a low or high quality product (at a low or high cost) and charge either a low or a high price. After the transaction, buyers do not learn which

---

[3]Also Huck et al. (2016a) study markets for experience goods. They focus on the effects of price regulation and price competition. Buyers have full information about the quality provided by each seller in the past. Since we do not deal with price regulation and have implemented private and not public information, our study is only marginally related to theirs.

quality they got.[4] With credence goods sellers can exploit clients in three ways, and DKS allow for all of them: undertreatment (providing a lower quality than needed), overtreatment (providing higher quality than needed), and overcharging (charging more than the good is worth). Reputation in DKS increases trade and decreases overcharging but does not decrease undertreatment and overtreatment, and has no effect on efficiency. Competition in DKS drives down prices and leads to maximal trade but has no effect on overtreatment, undertreatment, overcharging, and efficiency. When both competition and reputation are present, trade increases compared to the baseline but there is no further effect. DKS identify undertreatment as the main source of inefficiencies in their experiment.[5]

While all the previous studies use laboratory experiments, Schneider (2012), Rasch and Waibel (2013), and Balafoutas et al. (2013) test for inefficiencies in credence good markets directly in the field. Schneider (2012) takes a test vehicle to auto repair garages to check whether undertreatment, overtreatment, and overcharging occur and whether concern for reputation affects any of these (he signals either a motivation for a long lasting relation or a one-shot interaction). He finds that reputation does not improve outcomes. Rasch and Waibel (2013) complement the data from a field experiment similar to the one by Schneider (2012) with proxies for reputation and competition. According to their results, high competition decreases overcharging, while low concern for reputation increases it. Balafoutas et al. (2013) take a different perspective by looking at which customer characteristics lead to more overcharging in a field experiment with taxi rides.

Variations of the deception game (Gneezy, 2005) have already tested the efficacy of different

---

[4]This is the way DKS model credence goods. Namely, with credence goods, clients are not sure which quality they need, and after buying the good, they do not know which quality they got and whether they paid an appropriate price for it.

[5]Two studies build on the analysis of DKS. Kerschbamer et al. (2016) focus on the role of social preferences in explaining why credence goods markets with verifiability do not reach efficient outcomes. Mimra et al. (2013) extend DKS by investigating the role of public vs. private information about experts and compare the effect of fixed versus endogenously chosen competitive prices. Huck et al. (2016b) find that competition partially offsets the negative overtreatment effect of insurance in a credence good market.

remedies against cheating. For instance, in the monetary dimension, Peeters et al. (2008) look at rewards, while Sánches-Pagés and Vorsatz (2009) consider punishments, both after feedback about the quality of the message/advice. Angelova and Regner (2013) in contrast, focus on the role of voluntary payments offered to advisors not only after feedback about the transaction (bonus), but also before advice is given (upfront payment), and the combination of the two, the bonus and the upfront voluntary payment. The effect of information on truth-telling is tested by Ismayilov and Potters (2013) and Behnk et al. (2014): the former focus on the ex-ante disclosure of advisor's payoffs, while the latter study the effect of the (possible) disclosure of payoffs after the transaction. Van De Ven and Villeval (2015) measure the effect of scrutiny by a third party on the propensity of the sender to lie. Schotter (2003) summarizes a number of studies on advice, which, however, focus on situations without conflict of interest between advisors and clients.

To the best of our knowledge, ours is the first study that investigates the effect of voluntary payments, competition, and concern for reputation in a repeated deception game. The only other study that uses a partner design in a repeated deception game allowing for reputation building is Vanberg (2015). However, his focus is not on exploring the effect of reputation on truth-telling.

3   Experiment

We implemented an experimental deception game to study the effect of market forces (competition, reputation) and voluntary payments on the quality of the message (which we call advice). Subjects were randomly assigned a role of an advisor or client, which they kept throughout the entire experiment. The experiment consisted of 15 rounds. At the beginning of each round, only the advisors learned which state of the world was realized. State here is another word for the allocation of options to payoffs (in our case payoff pairs). Options were called A, B, C, and D. The payoff pairs were (10, 5); (5, 10); (5, 2); (5, 2) with the

payoff for the advisor listed first and that for the client – second. In the different states of the world, different payoff pairs were allocated to the same option. For instance, in one state of the world, option A gave 10 tokens to the advisor and 5 to the client; in another state, the same option yielded 5 to the advisor and 2 to the client. One possible state realization, as advisors saw it, is given in Table 3.

Table 1: A possible state realization

Option	Payoff for advisor	Payoff for client
A	10	5
B	5	10
C	5	2
D	5	2

Clients were informed about the possible payoff pairs, so that they were aware of the alignment of interests, as well as their own and the advisor's possible payoffs. However, clients were not informed what state of the world was realized, i.e., which payoff pair was assigned to which option. They had to choose one option, based solely on the advisor's recommendation. There were four possible recommendations the advisor could give. For example, recommendation 1 read: "With option A you will earn the most." Instead of showing the recommended option to the client, she was asked whether she wanted to follow the recommendation. If the answer was yes, the recommended option was implemented as her decision. If it was no, one of the other three options was randomly selected to be implemented as her decision. At the end of each round, both clients and advisors received feedback about which option was selected and their resulting payoffs. Advisors were also told whether the client followed the recommendation or not. Payoffs from the chosen option were added to their initial endowment of 2.5 tokens (paid in each round) to form the final payoff from the round. Two out of 15 rounds were randomly selected and paid out in the end of the experiment. One group of 10 subjects (5 advisors and 5 clients who only met subjects from the same group) qualified as one independent observation. When matching was exogenous (and this was treatment dependent), we chose a random stranger matching protocol. Within the 15

periods of the game, each advisor met each client 3 times, but we made sure that the same advisor did not meet the same client in two subsequent rounds. So far, this design closely follows the one in Angelova and Regner (2013), henceforth AR. Now, we will introduce the diverting features.

In all treatments advisors were first asked to pick a fee they would like to charge for their recommendation from the set of five possible fees: 0, 0.5, 1, 1.5, 2. After that they selected a recommendation based on the realized payoff table. Depending on the treatment, the fee was shown to either one particular client or to all clients within one matching group. In all treatments, the size (10 subjects), and the composition (5 clients and 5 advisors) of each matching group, as well as the matching protocol (endogenous or exogenous, if exogenous, then random, always clients with advisors) were common knowledge. An on-screen history box facilitated keeping track of one's own past interactions. It contained the period, the fee, the quality of the recommendation from the point of view of the client (good, medium, bad), and whether the client followed the advice. A final common feature of all treatments was that if a transaction did not take place, both the client and the advisor received their initial endowment (or outside option) of 2.5 ECU for this round.

In treatments *Base* and *Bon*, advisors and clients were randomly matched in pairs. Each advisor picked a fee, which was shown to her own client. If the client was willing to pay the fee, she would receive the advice, otherwise both the client and the advisor would earn their 2.5 tokens from this round. If the client got a recommendation, she would decide whether she wanted to follow it. After that, everyone received feedback about own earnings. Additionally, in *Bon*, clients would be able to offer a bonus to the advisor up to the amount of the client's total earnings in this round. The history of the round was summarized in the info box which in treatment *Bon* additionally listed the bonus paid/received.

In treatments *Comp* and *CompBon*, clients were informed about all fees in a random order.[6]

---

[6]Fees would be displayed in one row, but subjects knew that their order was determined by chance in each round, such that it was not possible to detect a particular advisor based on the position of her fee in

Based on the fees clients chose their preferred advisor for each interaction. While each client would choose maximally one advisor, each advisor could be selected by a number of clients between zero and five. It was common knowledge that independently of the number of interactions, each advisor would charge the same fee and send the same recommendation to all her clients.[7] The total payoff of an advisor from a given period is equal to the sum of payoffs from all interactions in that period. At worst, an advisor would not be chosen for an interaction; her period-payoff in this case would be equal to her initial endowment. This initial endowment serves as the safe outside option that advisors have independently of the treatment. The history box for advisors was augmented with two additional pieces of information: the number of clients served and how many of them followed the advice.

In treatments *ID* and *IDBon*, each advisor received a unique identification number. The history box was augmented by a column containing that identification number. It enabled clients to identify advisors they already interacted with. Clients did not receive any information about the quality of advice provided to other clients. Hence, the only information clients had about a particular advisor was based on their previous experience with that advisor.

Finally, in treatments with competition and identifiability (*IDComp, IDCompBon*), clients would pick an advisor based both on fees for advice and their previous experience with a particular advisor.

Table 2: Treatments

	Without competition	With competition
Baseline	*Base*	*Comp*
Identifiability	*ID*	*CompID*
Bonus	*Bon*	*CompBon*
Identifiability & Bonus	*IDBon*	*CompIDBon*

See Table 2 for an overview of the treatments. We ran two sessions per treatment. Each

the row.

[7]Since clients are not identifiable to advisors, it does not make sense to allow for discriminatory advice.

session consisted of 30 subjects, such that per treatment there were 60 subjects (30 advisors, 30 clients).

Table 3 summarizes the main design features of our study compared to the three experiments that are most related to ours: DKS, HLT, and AR. The first two rows show percent changes in principal's[8] pay when entering the game relative to the outside option (i.e., the safe pay principals receive when they choose not to interact). The increase or decrease in pay relative to the outside option is a measure for the attractiveness to enter the game. To compute changes in pay, we took into account the outside options, the payoffs of principals when they enter the game and get 'fully exploited'[9], and the payoffs of principals when they enter the game and receive the maximal possible payoff.

The outside option in DKS is 1.6, in HLT it is 20, in AR there is no outside option, and in our study it is 2.5. In fact, in our study, if players don't interact, they get only 2.5 each, while if they do interact, they get their earnings from the game on top of those 2.5. When maximally exploited, the buyer in DKS obtains the negative payoff of –1 (the buyer's payoff function is 10–price, and prices are discrete numbers from 1 to 11). In HLT, the trustor receives 5 when she trusts but her trust is not honored. In our study the client is maximally exploited when the most favorable option for the advisor is implemented and the client is charged the maximal fee, hence 5 minus 2 plus 2.5. The maximal payoffs for principals when they enter the game are 9 in DKS, 30 in HLT, and 12.5 in our study. Thus, considering, e.g., the client in our study, when fully exploited, she gets 5.5 and the outside option is 2.5. Hence, her minimum pay relative to the outside option is $\frac{5.5-2.5}{2.5} \cdot 100 = 120\%$. The maximum pay relative to the outside option is $\frac{12.5-2.5}{2.5} \cdot 100 = 400\%$. The rest of the table is self-explanatory. In the following we will discuss the differences between our study and the other three.

In comparison to DKS it is far more attractive to enter the game in our study: even if fully

---

[8]I.e. buyers in DKS, trustors in HLT, and clients in AR and our study.

[9]Agents get the maximal possible payoff for themselves at the expense of principals.

Table 3: Design features of DKS, HLT, AR, and our study

	DKS	HLT	AR	Our study
Principals' min. pay relative to outside option*	−162.5%	−75%	–	+120%
Principals' max. pay relative to outside option*	+462.5%	+50%	–	+400%
Competition in fees only	yes	no	no	yes
Competition in reputations only	no	yes	no	no
Competition in reputations and fees	yes	no	no	yes
Cheating detectable	no	yes	yes	yes
Undertreatment**	yes	yes	yes	yes
Overtreatment	yes	no	no	no
Overcharging	yes	no	no	yes
Bonus	no	no	yes	yes
Periods	16	30	15	15
Size of matching groups	8	8	10	10
Reputation via private info	yes	yes	no	yes
Reputation via public info	no	yes	no	no

Notes: * Principals' (buyers in DKS, trustors in HLT, clients in AR/our study) minimum pay is their payoff when fully exploited in the interaction with the agent. Likewise, maximum pay is the highest possible payoff principals can make. ** In our game joint payoffs for advisor and client are constant (unless a (5, 2) – recommendation is followed). Nevertheless, not implementing the best option for the client implies an efficiency loss as abuse of the commission system via advice steering results in low-performing investments. See Inderst and Ottaviani (2012) for a formal model and Angelova and Regner (2013) for details.

exploited, payoffs of the client *increase* by 120% when she enters the game. This is not the case in DKS where buyer's earnings *decrease* by 162.5%. A further difference to DKS is that in our study advice is modeled as an experience and not as a credence good. That is, clients know what quality they need and learn at the end of each period whether they got this quality (in DKS cheating is not detectable). Hence, our design rules out overtreatment but allows for undertreatment and overcharging. Giving truthful advice corresponds to a high quality good. Telling a lie is equivalent to undertreatment and charging a positive fee for a lie is equivalent to overcharging. While DKS also analyze the role of institutions (liability,

verifiability), we allow for voluntary payments to study the role of reciprocity.

Also in HLT entering the game is much less attractive than in our study: if exploited, the trustor's earnings *decrease* by 75%. Moreover, in HLT the trustor can maximally increase earnings by 50% when entering the game. In our study, the increase in pay in that case amounts to 400%. Hence, the trustor in HLT has much more power than our client, since the trustor can credibly threat not to enter the game. A further difference to HLT is that competition in our study takes place along two dimensions: clients choose an advisor based both on fee for advice and quality of advice given in the past. We lack a treatment, in which advisors are selected only given the information about their past behavior, and HLT lack a treatment in which advisors are selected only given their fees.

Finally, the differences between AR and our study are that AR are missing an outside option, a treatment with competition, and a treatment where advisors are identifiable.

3.1   Procedures

The experiment was conducted with students from the University of Jena. They were invited to the laboratory of the Max Planck Institute of Economics using the online recruitment system for economic experiments ORSEE (Greiner, 2004). The experiments were computer-based, using z-Tree (Fischbacher, 2007). Subjects earned 19.07 Euros on average and spent between 90 and 120 minutes (30 minutes of which on the instructive part) in the laboratory.

Upon arrival in the laboratory, subjects were randomly assigned to a cubicle, where they individually read the instructions (see appendix). During the experiment, eye contact was not possible. Although participants saw each other at the entrance of the lab, there was no way for them to guess with whom of the other 29 students they would be matched later on. All subjects had participated in at least one experiment before.

3.2 Behavioral Predictions

We begin with the analysis of the stage game.[10] Then, we relax the standard self-interest assumption in order to accommodate empirical evidence of honest behavior in similar settings. Thereafter, we consider repeated game aspects in the identifiability treatments. In the end, we state our hypotheses.

The stage game is a dynamic game, in which the advisor's strategy consists of two actions: choosing a fee for advice and choosing a recommendation. The client's strategy also consists of two actions: deciding whether to pay the fee and, if yes, whether to follow the advice. The natural sequence of decisions is the following. The advisor first chooses a fee, the client is then informed about the fee and decides whether to pay it. If not, the game ends and both players receive a payoff of 2.5 each. If yes, then the advisor sends a recommendation, and finally the client decides whether to follow that recommendation.[11] The subgame, in which the advisor recommends an option (A, B, C, or D) and the client can decide between following (F) and not following (NF) is depicted in Table 4.

Table 4: Subgame played after the advisor chooses a fee and the client agrees to pay it.

Advisor \ Client	F	NF
recommend A	(10; 5)	(5; 4.7)
recommend B	(5; 10)	(6.7; 3)
recommend C	(5; 2)	(6.7; 5.7)
recommend D	(5; 2)	(6.7; 5.7)

The subgame has three Nash equilibria in pure strategies, (A, F), (C, NF) and (D, NF),

---

[10]The stage game is the game played in each of the 15 periods. Clearly, in the treatments without identifiability, subjects play a series of 15 one-shot games. In the treatments with identifiability, meeting the same opponent again turns the one-shot game into a repeated game.

[11]In the experiment we used the strategy method for advisors, i.e. we asked them to choose the fee and the recommendation simultaneously and without knowing whether the client decided to pay the fee or not.

as well as an infinite number of Nash equilibria in mixed strategies. It is always reached in equilibrium, because the client always chooses to pay the fee and enter the subgame instead of getting the outside option of 2.5. The advisor always charges the highest possible fee. For all subgame perfect equilibria in the stage game one can easily find supporting beliefs and the corresponding perfect Bayesian equilibria.

Notice that in both subgame perfect equilibria in pure strategies, advisors are predicted to lie (i.e., recommend anything but option B) to their clients. However, similar experiments (e.g., Gneezy, 2005; Cai and Wang, 2006; Sánchez-Pagés and Vorsatz, 2007; Hurkens and Kartik, 2009; Erat and Gneezy, 2012; Danilov et al., 2012; Angelova and Regner, 2013) frequently find less lying than predicted, a behavioral pattern known as "overcommunication phenomenon". Social preferences are the widely accepted explanation for this behavior. More specifically, people may be lying less than standard theory predicts because of an aversion to cheat, an aversion to feeling guilty because of disappointing the counterpart, or out of reciprocity/fairness concerns. Treating the precise nature of social preferences as a black box, we assume that whether advice is given honestly depends on two factors: advisors' *pro-sociality* ϕ and the *situational context* λ.

It may be easiest to interpret *pro-sociality* $(0 \leq \phi \leq 1)$ as an individual's self-image about being honest.[12] The higher ϕ, the likelier it is that honest advice is given. If $\phi = 1$ the advisor always gives honest advice. This would correspond to an advisor with an immaculate self-image about being honest. In the extreme case of a purely selfish individual $\phi = 0$ and honest advice is never given.

Moreover, the behavior of advisors may be affected by the *context of the situation* $(\lambda > 0)$. In a transparent baseline situation, $\lambda = 1$. External factors that result in a higher chance to give truthful advice are expressed by $\lambda > 1$. Such an increased tendency to advice

---

[12]Models of self-image concerns provide a theoretical basis for honest behavior, see the literature on cognitive dissonance (Festinger, 1957), identity (Akerlof and Kranton, 2000), self-concept maintenance (Mazar et al., 2008), self-signaling (Bénabou and Tirole, 2011).

honestly could depend on a variety of external factors.[13] In our experimental design mutual opportunities to reciprocate serve as the instrument to create a favorable situational context when the option to give a bonus exists. Hence, our design varies the situational context, in the dimension of reciprocity, that is, $\lambda = 1$ in treatments without bonus and $\lambda > 1$ in treatments with bonus. We allow advisors' reaction to a change of the situational context to be individually heterogenous but assume that their reaction to a change of λ is equally distributed over ϕ and treatments. Thus, on average, a change of λ results in a change of behavior.

Combining the effects of pro-sociality and situational context, we express an advisor's tendency to lie in the stage game by $l = \frac{(1-\phi)}{\lambda}$.[14] When $\phi = 1$, an individual will always give truthful advice, irrespectively of the situational context λ. For $\phi < 1$, a favorable situational context ($\lambda > 1$) decreases the tendency to lie in comparison to the baseline ($\lambda = 1$). The binary variable $h(l)$ expresses whether an advisor is honest (1) or not (0). Its value depends on whether an advisor's tendency to lie is above a threshold level $0 < l^* < 1$:

$$h(l) = \begin{cases} 0, & \text{if } l > l^* \\ 1, & \text{otherwise} \end{cases}$$

Finally, *reputation concerns* may affect the decision to advice truthfully or not in the finally

---

[13]For instance, pre-play communication between the agents, in particular making a promise, has been found to reduce cheating (Charness and Dufwenberg, 2006; Beck et al., 2013). The relative monetary costs of lying seem to matter (Erat and Gneezy, 2012). Opportunities to reciprocate, especially if they are mutual, lead to more truthful advice in Angelova and Regner (2013). In related settings (charitable giving, public goods games), reduced anonymity leads to more pro-social behavior. In contrast, a non-transparent situation that provides moral excuses for dishonest behavior would result in $\lambda < 1$ (see the moral wiggle room literature following Dana et al., 2007).

[14]We assume this functional form to focus on the two key effects: a higher degree of pro-sociality and a favorable situational context result in less lies. We do not rule out alternative specifications.

repeated game with identifiability.[15] Following the standard set up of a Kreps et al. (1982) reputation model, we distinguish two types of dishonest advisors ($h(l) = 0$). Strategic ones maximize their profit over the repeated game. They invest in reputation by imitating honest advisors as long as their future benefits from this reputation warrant the investment. In contrast, myopic dishonest advisors do not look beyond the current round. They never give truthful advice as they always maximize their stage game profits. In accordance with the standard reputation model we assume that honest, strategic dishonest and myopic dishonest advisors are represented at substantial levels and that pro-sociality is randomly distributed across treatments. Thus, when reputation concerns matter strategic dishonest advisors have an incentive to imitate honest advisors by giving truthful advice as long as their future payoffs make it worthwhile, while myopic ones behave like in the stage game and do not give truthful advice.

In treatments *ID* and *IDBon*, in addition to playing the stage game equilibria in each period, it is possible to construct equilibrium strategies, in which sending a truthful recommendation and following the advice the first time an advisor–client pair interacts can be part of a subgame perfect equilibrium.[16] In the identifiability treatments under competition, *CompID* and *CompIDBon*, clients can credibly threat to stop interacting with advisors who do not tell the truth. It is straight forward to show that the following strategies constitute a subgame perfect equilibrium in the finally repeated game with competition and identifiability: advisors charge a fee of zero and send a truthful recommendation in all periods but the last one; clients pay the fee, and follow the recommendation in all periods; if the advisor sends a truthful recommendation, the client selects the same advisor in the subsequent period, if not, the client switches to a different advisor. In the last period, the equilibrium (A, F) is played in

---

[15]The pro-sociality of subjects is assumed to be constant over the course of the game.

[16]Recall, that in this treatment an advisor will meet a client 3 times during the 15 periods of the game. Hence, the client can use the credible threat of enforcing one of the equilibria (C, NF) or (D, NF) in the second and third interaction, if the advisor does not give truthful advice in the first. If the advisor gives truthful advice in the first interaction, then the client will enforce the more beneficial equilibrium for the advisor (A, F) in the last two interactions.

the subgame.[17]

3.2.1 Treatments without competition

We begin with treatment *Base*. Advisors behave according to $h(l(\phi, \lambda = 1))$.

H1: In Base the rate of truthful advice is significantly greater than zero.

In treatment *ID*, reputation concerns may motivate advisors. Since they apply only to the first third of the experiment, we expect a relatively weak effect on truth-telling.

H2: The rate of truthful advice in ID is greater than in Base.

In treatment *Bon*, the possibility to give a bonus provides a one-sided opportunity (for the client) to reciprocate. It can be seen as a situational context that induces honest behavior via the channel of reciprocity ($\lambda_{Bon} > 1$). Pro-social clients may reciprocate receiving truthful advice by giving a bonus. Hence, an advisor's tendency to give truthful advice increases, the more he expects to meet a reciprocating client.

Angelova and Regner (2013) find a sustainable positive effect of the combination of upfront voluntary payment and bonus afterwards. However, they also report a positive effect but a decay over time if the bonus stands alone. In our design the bonus is the only voluntary component as the upfront fee is charged by the advisor and not voluntarily offered by the client. Hence, we cautiously expect a tendency among advisors to be more truthful.

H3: The rate of truthful advice in Bon is greater than in Base.

---

[17]If all advisors follow this strategy, every advisor will interact with, on average, one client for 15 periods. If an advisor unilaterally deviates, she will lose her client(s) and earn only the outside option until the end of the game, which is always less profitable than sticking to the equilibrium strategy. For the client it is also never profitable to unilaterally deviate, since in every period she gets the highest possible payoff; in the last period, an equilibrium is played.

In treatment *IDBon*, honest advisors can be identified. A reciprocal relationship between an advisor and a client (truthful advice, bonus paid) in one interaction can now extend to subsequent meetings. The advisor has a chance to reciprocate the paid bonus by giving truthful advice when they meet again. In this way mutual opportunities to reciprocate arise. As a consequence, the negative effect of the situational context on the tendency to lie is amplified and we expect a higher tendency to advise truthfully than in the other treatments without competition ($\lambda_{IDBon} > \lambda_{Bon} > 1$).

H4: The rate of truthful advice in IDBon is greater than in Base, ID or Bon.

3.2.2 Treatments with competition

Treatment *Comp* serves as a baseline treatment within the set of the treatments with competition.

In *CompBon*, pro-social clients can reciprocate truthful advice within one interaction by giving a bonus. Thus, the situational context is more favorable than in *Comp* ($\lambda_{CompBon} > 1$). Moreover, in the competitive environment advisors are able to advise more than one client. As these clients might reciprocate truthful advice with a bonus, the incentive for telling the truth is amplified. Thus, we expect an increased tendency to advice truthfully.

H5: The rate of truthful advice in CompBon is greater than in Comp.

If advisors are recognizable, strategic advisors, who would otherwise lie, are exptected to be honest in every period except the last one. Therefore, we expect increased truth-telling and a drop in the rate of truthful advice in period 15. Myopic advisors will play one of the stage game equilibria in every period.

H6: The rate of truthful advice in CompID is greater than in Comp.

In *CompIDBon* advisors are identifiable when competing for clients and clients can pay a

bonus. Thus, mutual opportunities to reciprocate exist between advisor and client, because they could meet again in a subsequent period. In the extreme, an advisor-client relationship lasts for all 15 periods. This environment fosters reciprocity concerns among advisors ($\lambda_{CompIDBon} > \lambda_{CompBon} > 1$). We expect that this reciprocity effect further strengthens the positive effect of reputation concerns on the tendency to give truthful advice.

H7: The rate of truthful advice in CompIDBon is greater than in Comp, CompBon or CompID.

4 Results

In this section we present the choices of advisors and clients, and test for the effects of our treatments on their behavior. Finally, we take a closer look at the dynamics in the competition treatments.

4.1 Choices of advisors and clients

Table 5 provides an overview of the average choices of advisors and clients. Advisors post a fee of around 1, on average. Competition significantly lowers mean posted fees from 1.13 (in treatments without competition) to 0.74 (in treatments with competition), Wilcoxon-Mann-Whitney-Test, $p < .01$. Mean posted fees do not statistically differ from mean accepted fees in all treatments with competition but *Comp*, where the mean accepted fee is higher than the mean posted fee, Wilcoxon signed-rank test, $p < .07$.

In *Base* clients took the outside option in 9% of all transactions. This is the highest observed rate. In treatments with competition essentially no client decided to take the outside option.

The rate of truthful advice in *Base* is 27%, the lowest rate of all treatments. It is highest

Table 5: Descriptive statistics

	Base	Bon	ID	IDBon	Comp	CompBon	CompID	CompIDBon
Average posted fee	0.98	1.08	1.07	1.39	0.93	0.83	0.65	0.59
Average accepted fee	–	–	–	–	1.01	0.79	0.52	0.55
Outside option	9%	5%	7%	6%	2%	1%	1%	0%
Truthful advice	27%	33%	32%	56%	44%	62%	74%	75%
Follow	77%	88%	79%	91%	92%	95%	98%	98%
Truthful advice followed	19%	28%	25%	52%	48%	67%	81%	84%
Bonus paid	–	34%	–	47%	–	48%	–	68%
Average bonus	–	1.34	–	1.67	–	1.64	–	1.71

with identifiability of advisors and competition in the same treatment (74% in *CompID* and 75% in *CompIDBon*). Figure 1 depicts the rates of truthful advice over time. The left panel shows treatments without competition, the right panel those with competition. In all treatments without competition we observe a downward tendency over time. In the competition treatments behavior is stable until the last three periods.
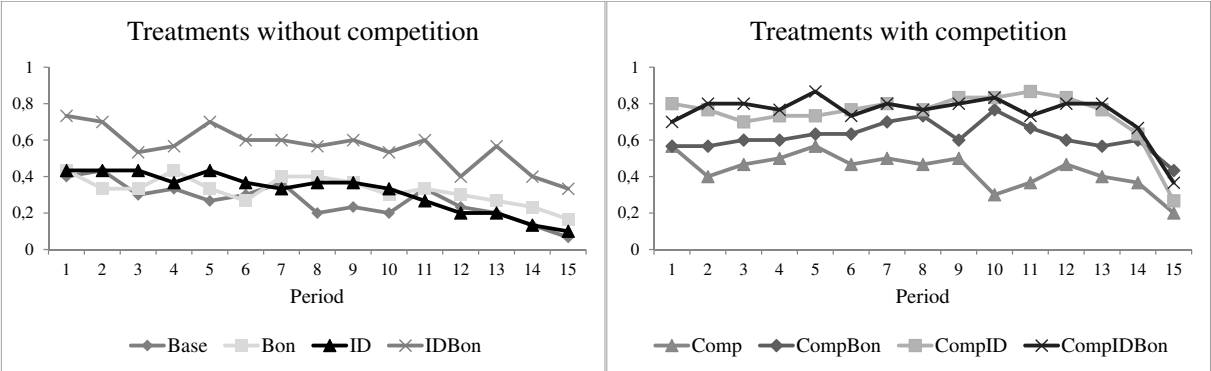

Figure 1: Truthful advice over time by treatment

In *Base* clients decide to follow in 77% of all interactions with an advisor. This is the lowest rate we observe across treatments. In *IDBon* as well as in all the competition treatments the follow rate is above 90%. In 34% of all *Bon* interactions clients paid a bonus and if they did, on average, this bonus was 1.34. In contrast, 68% of all *CompIDBon* interactions included a bonus and the average bonus was 1.71. Figure 2 illustrates the rate of bonus payments and their average size over time. We observe a substantial end-game effect in periods 14 and 15. The average bonus in period 15 drops down to 0 in *Bon*, 0.2 in *IDBon* and 0.3 in the treatments with competition. Before that the average bonus appears stable in treatments *CompBon* and *CompIDBon*. The large majority of clients paid a bonus at least once: over 90% in *Bon*, *IDBon*, and *CompIDBon*, and 77% in *CompBon*. Moreover, many clients frequently paid a bonus. The percentage of clients who paid a bonus more than half of the time was 60% in *IDBon* and *CompBon*, 80% in *CompIDBon*, and 27% in *Bon*.
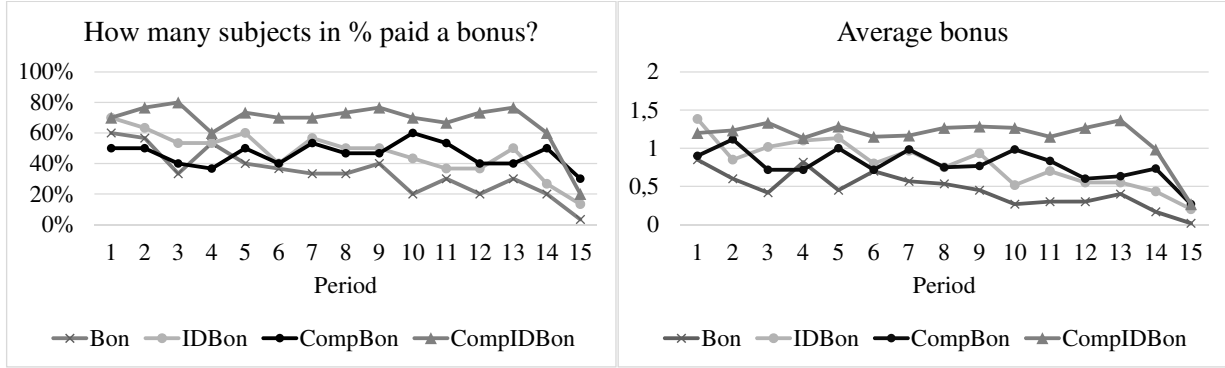
Figure 2: Percentage of bonus-payers and average bonus over time

4.2 Treatment comparisons

In order to test for treatment effects we set up a panel that contains all 3,600 interactions between advisors and clients. Table 6 reports the results of two logit mixed effects regressions with random terms associated with matching groups and advisors.[18] The dependent variable is whether truthful advice has been given (1) or not (0). Explanatory variables are the fee posted by the advisor and dummy variables for the treatments. In order to control for the apparent negative time trend we include a dummy for the period and a dummy for the last period.

Specification 1 compares the treatments *Bon*, *ID* and *IDBon* to *Base*. We find a positive correlation between the posted fee and truthful advice (significant at the 1%-level). Neither the dummy for *Bon* nor the dummy for *ID* are significant. The dummy for *IDBon* is positive and highly significant. The coefficient of *IDBon* (2.39) is greater than the one of *ID* (0.81) or *Bon* (0.71) ($p < 0.05$). The period dummy as well as the dummy for the last period are negative and significant at the 1%-level.

Specification 2 compares the treatments *CompBon*, *CompID* and *CompIDBon* to *Comp*. Again, the posted fee and truthful advice are positively correlated (1%-level). All treatment

---

[18] All reported results are robust to using standard random-effects logit models.

dummies are significant at the 1%-level. While the coefficient of *CompIDBon* (2.79) is greater than the one of *CompBon* (1.47) ($p < 0.05$), it is not significantly greater than the one of *CompID* (2.58). We do not find evidence for a negative time trend, only for a drop in the last period (significant at the 1%-level).

Table 6: Determinants of truthful advice				
	1: Without competition		2: With competition	
Posted fee	0.88***	(0.1)	1.10***	(0.1)
Bonus	0.71	(0.7)	1.47***	(0.6)
Identifiability	0.81	(0.7)	2.58***	(0.6)
Identifiability + Bonus	2.39***	(0.7)	2.79***	(0.6)
Period	-0.16***	(0.02)	-0.00090	(0.02)
Last Period	-1.01***	(0.4)	-2.68***	(0.3)
Constant	-1.77***	(0.6)	-1.35***	(0.5)
Observations	1800		1800	

Standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Logit mixed effects regressions with random terms associated with matching groups

and advisors; dependent variable: truthful advice (1 if given, 0 if not)

The positive but non-significant effects of the *ID* and of the *Bon* dummy on truthful advice reject hypotheses 2 and 3. However, we do find a positive and significant effect of *IDBon*, thus supporting hypothesis 4. With competition all treatment dummies are positive and significant. Hence, hypotheses 5 and 6 are supported. However, the *CompIDBon* treatment does not result in a significant increase of the rate of truthful advice in comparison to *CompID*. Hypothesis 7 is only partly supported.

To summarize, when there is no competition we find the hypothesized positive effect of a bonus only in combination with identifiability. It seems that additional interaction between client and advisor is necessary, that is, over periods as they know they meet again and not only within a period. Under competition, as expected, identifiability as well as the bonus

increase the rate of truthful advice. However, we do not find an additional positive effect when identifiability and bonus are combined.

We proceed with a test of treatment effects on the clients' decision to follow advice. Table 7 reports a set of logit mixed effects regressions. The dependent variable is whether the client followed the advice (1) or not (0). Hence, observations are dropped when the client decided against taking advice in the first place. Explanatory variables are the fee charged by the advisor, dummy variables for the treatments as well as the period and a last period dummy. Specification 1 presents results for the treatments without competition. Specification 2 adds a dummy whether in the previous period the client had a good experience, that is, whether she followed good advice. Specifications 3 and 4 show respective results for the competition treatments.

In the treatments without competition, both specifications yield similar results. The follow rate is positively correlated with the posted fee (significant at the 5% level). In *Bon* and *IDBon* the follow-rate is higher than in *Base* (significant at the 5% and 1% level, respectively). While the coefficients of *IDBon* are greater than the ones of *ID* ($p < 0.01$), they are not significantly greater than the ones of *Bon*. The more experienced a client gets, the less likely she is to follow the advice (the coefficient for 'period' is negative and significant at the 5% level at least). Whether the client made a good experience with her interaction in the previous period does not affect her decision to follow the advice in the current period.

In the treatments with competition, again the higher the fee the higher the follow-rate (significant at the 1% level). In *CompBon* the follow-rate does not differ from that in *Comp*. Clients in *CompID* are slightly more likely to follow the advice than in *Comp* (significant at the 10% level). However, in *CompIDBon* clients follow the advice significantly more often than in *Comp* (significant at 5%). The coefficients of *CompIDBon* are not greater than the ones of *CompID* or *CompBon*, though. In contrast to the treatments without competition, the more experienced clients get, the more likely they are to follow advice (significant at 1%). In the last period, the follow rate drops significantly ($p < 10\%$). Finally, having received

	Without competition		With competition	
	(1)	(2)	(3)	(4)
Posted fee	0.28**	0.27**	2.63***	2.56***
	(0.1)	(0.1)	(0.5)	(0.5)
Bonus	0.80**	0.87**	1.33	1.60
	(0.3)	(0.4)	(1.0)	(1.0)
Identifiability	-0.010	0.063	2.03*	1.73*
	(0.3)	(0.3)	(1.0)	(1.0)
Identifiability + Bonus	1.19***	1.21***	2.76**	2.81**
	(0.4)	(0.4)	(1.1)	(1.1)
Period	-0.073***	-0.057**	0.16***	0.15***
	(0.02)	(0.02)	(0.04)	(0.05)
Last Period	0.14	0.083	-1.27*	-1.36*
	(0.3)	(0.3)	(0.7)	(0.7)
L.GoodAdviceFollowed		0.033		1.36***
		(0.2)		(0.4)
Constant	1.83***	1.64***	1.23	0.62
	(0.3)	(0.3)	(0.8)	(0.8)
Observations	1678	1469	1778	1641

Table 7: Determinants of the decision to follow

Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Logit mixed effects regressions with random terms associated with matching groups and clients; dependent variable: advice followed (1 if yes, 0 if no)

good advice in the previous period appears highly important as it is positively correlated to the decision to follow (significant at the 1% level).

To sum up, fees are positively related to the follow rates independently of the treatment. Clients are more likely to follow the advice in those treatments, in which advisors are identifiable and bonus payments are possible at the same time. Without competition, the follow-rates decrease over time, while with competition, they increase. Positive experience with truthful advice from the previous period increases the probability to follow the advice in the current period only in the treatments with competition.

4.3   Dynamics in the competition treatments

Table 8: Market shares in the competition treatments

	Comp	CompBon	CompID	CompIDBon
Highest market share	25%	27%	53%	49%
2nd highest market share	21%	24%	20%	25%
3rd highest market share	20%	20%	13%	16%
2nd lowest market share	17%	16%	9%	7%
Lowest market share	15%	12%	5%	3%

In the following, we compare advisors' market shares across the competition treatments, see Table 8. We calculated the market share (i.e. the number of clients served divided by the total number of clients in the matching group) for every advisor in every period. Next, we ranked the market shares from highest to lowest for each period and each matching group, such that for every period-matching group combination five market share categories (highest to lowest) result. Finally, we averaged the entries in each category over the periods and matching groups.

In the treatments where advisors are not identifiable, market shares are quite equal. This is not surprising because by design in these treatments clients can choose an advisor only based on posted fees, and identifying the possibly honest advisor from the previous period is unlikely. In contrast, when advisors are identifiable, market shares become very unequal. For instance, in *CompID* market shares range from 53% (being the highest) to 5% (being the lowest). Obviously, some advisors manage to attract and keep the majority of clients. What is the key to a large market share in the *ID* treatments?

For every matching group in the *ID* treatments we identified the advisor with the highest market share and analyzed her strategy regarding posted fees and advice quality. It turns out that two things are crucial for a large market share: first, to be selected to advise as many clients as possible already in period 1, and second, to keep advising truthfully.

Even if advisors in the *ID* treatments intend to give truthful advice, failing to attract clients in the beginning, puts them at risk of an empty store for the rest of the game, given that competitors remain honest, and thus their clients have no reason to switch. In early periods, the only way to attract clients is by choosing the 'right' fee, that is, the fee that will be selected by most clients. Table 9 gives an overview of the percentage of clients who chose the lowest or the highest fee, both for period 1 and the entire game. Looking at period 1, in the *ID* treatments most clients (50–57%) pick the lowest posted fee(s). Indeed, in both *ID* treatments, in period 1, the average accepted fee is significantly below the average posted fee.[19] Hence, the secret of attracting clients in the *ID* treatments seems to be posting a low fee in period 1. Notice, however, that for the *ID* treatments the column referring to the entire game is not very informative because a client may be choosing a high fee not because she likes expensive fees but because this may be the only way to remain affiliated with the same advisor.

---

[19]In *CompID*, in period 1 the mean posted fee is 1.18, the mean accepted fee is .85; both differ significantly at the 3% level (Wilcoxon signed-rank test). In *CompIDBon*, in period 1 the mean posted fee is .87, the mean accepted fee is .72; both differ significantly at the 5% level (Wilcoxon signed-rank test).

Table 9: Do clients select the lowest fee?

	Percentage of clients who selected lowest vs. highest fee	
	in period 1	in all periods
Comp	37% vs. 17%	28% vs. 32%
CompBon	40% vs. 17%	43% vs. 26%
CompID	57% vs. 20%	54% vs. 24%
CompIDBon	50% vs. 17%	44% vs. 28%

What is the key to keeping those clients? It is to give truthful advice period after period. Successful advisors maintain long term interactions with the same two or three (seldom four or five) clients. Most advisors give truthful advice until the very last period of interaction. In contrast, dishonest advice in one period leads to an immediate loss of clients, even if advisors have been honest for a number of periods before that. Similarly, most clients switch to a different advisor if their (although truthful) advisor starts increasing the fee (too much) compared to the own fee in previous periods and the fees of the other competitors in the same period.

In treatments without identification (*Comp* and *CompBon*), the only tool to attract clients is the posted fee. But what is the 'right' fee? In these treatments, in period 1 there is a tendency to choose the lowest posted fee (37–40%) and not the highest posted fee (17%), see Table 9. However, the remaining around 43%–46% of clients select any posted fee in between. Also, there is no significant difference between mean posted and mean accepted fees in period 1. Considering the entire game, the percentage of clients willing to pay the highest posted fee is higher than in period 1: 32% in *Comp* and 26% in *CompBon*. Indeed, in *Comp*, the mean accepted fee is significantly above the mean posted fee, when considering the entire game. It seems that in later periods in *Comp*, clients start to believe that higher fees give them a higher chance to receive truthful advice. The regressions in Table 6 show indeed a positive correlation between posted fees and truthful advice. At least in treatment *Comp* clients realize this relationship. However, a large amount of clients chooses anything between the highest and the lowest fee. So all in all, there is no consensus among clients

which is the most attractive fee and, hence, advisors cannot really employ any clear strategy to attract clients.

4.4 Discussion

Keeping in mind that the experimental paradigms are different across DKS, HLT and our study, we compare behavior in the baseline treatments (i.e. *B/N* in DKS vs. *control* in HLT vs. *Base* in our study), as well as behavior when both competition and reputation concerns apply (i.e. *CR/N* in DKS vs. *pi-c* in HLT vs. *CompID* in our study). In the baseline treatments, HLT report that trust is being honored in 28% of the cases, in our study the rate of truthful advice is 27% and in DKS sellers provide appropriate quality in 47% of the cases.[20] With competition and reputation, HLT document an honor rate of 92%, we find truthful advice in 74% of the cases, while DKS report appropriate quality in 36% of the cases. Cheating being perfectly detectable in HLT and our study seems to be the most probable explanation for this difference to DKS. Note that Rubin and Sheremeta (2015) find a similar result in the context of reciprocity: when random shocks distort agent effort, wages and effort levels in their gift exchange game drop.

The use of the bonus is relatively prevalent: across treatments between 34% and 68% of all transactions include a positive bonus. In treatments *IDBon*, *CompBon* and *CompIDBon* more than half of all clients paid a bonus in more than half of their transactions. For the bonus option to have an effect on the truthfulness of advice, multiple opportunities to reciprocate appear to be necessary. We find a significant increase of the rate of truthful advice only if clients interact with advisors not just within one period (*Bon*) but several times over the course of the game (*IDBon*, *CompIDBon*) or several clients can reciprocate within one period (*CompBon*). The condition of opportunities to reciprocate being mutual

---

[20]Appropriate quality in DKS is 100% minus the rate of undertreatment. In order to provide an adequate comparison with our study where overtreatment is ruled out, we just consider the rate of undertreatment in DKS.

is consistent with the findings of Angelova and Regner (2013).

Given that adding a bonus or identifiability to competition increases the rate of truthful advice significantly, it appears puzzling why adding both does not lead to a further increase in the rate of truthful advice. A possible explanation is a ceiling effect. Since in *CompID* the rate of truthful advice is already 74%, possibly only myopic selfish advisors remain and thus adding the opportunity to pay a bonus does not induce additional advisors to switch to giving truthful advice.

Finally, our setting differs from the trust game in HLT and the game in DKS in that cheating in our deception game is more pronounced. By choosing not to recommend the best option for the client, an advisor in our game explicitly tells a lie. If there is a moral cost to lying, subjects in our study could be expected to behave pro-socially more often than subjects in the other studies. Moreover, they could be more sensitive to incentives that lead them to lie less. However, the small differences between the results of HLT and our study do not indicate that moral costs of lying are substantial.

5 Conclusions

In a deception game (Gneezy, 2005), we study experimentally possible remedies against moral hazard, i.e. misleading advice given to clients. We introduce competition among advisors, the possibility for them to build a reputation, and a channel through which clients can reciprocate if they got truthful advice: a voluntary bonus paid after feedback about advice quality.

Without competition, mutual opportunities to reciprocate lead to the provision of significantly more truthful recommendations. They exist when the option to give a bonus is coupled with advisor identifiability allowing interactions over time not only within one period. In the competition treatments, the bonus or identifiability significantly increase the

rate of truthful advice. However, we find no further increase when competition, bonus and identifiability are combined, possibly due to a ceiling effect.

Comparing our results to related studies of experience/credence goods, the combination of competition and reputation concerns also leads to the lowest rate of opportunistic behavior in Huck et al. (2012), while it has no effect in Dulleck et al. (2011). It seems that cheating being perfectly detectable – a common feature of our study and Huck et al. (2012) – is a pre-condition for a positive effect of reputation and competition.

Being able to rely on market forces like competition and reputation in order to foster efficiency seems reassuring. However, in real life settings implementing competitive environments and reputation mechanisms may not always be straightforward. For instance, the effect of 'private' reputation (based on clients' own experiences) requires repeated interaction with the same advisor. Without anticipating a potentially long-lasting relationship to a client, advisors may not be willing to invest in reputation. Moreover, first-time clients are precluded from accessing information about the advisor altogether. Hence, if client-advisor relationships are relatively short-term, incentives to advise truthfully do not really kick in and 'private' reputation does not appear to be a useful instrument.

Thus, our result of a bonus effect points at a possible safeguard against opportunistic behavior in market environments where asymmetric information and conflicting interests would otherwise lead to inefficient outcomes. The voluntary component activates reciprocal concerns and, combined with the possibility of reputation building or having multiple clients, decreases cheating by advisors, increases the follow rate of clients, and leads to more efficiency.

What are potential applications of our 'bonus' in reality? Beyond the literal interpretation of an actual monetary payment of a content client to a truthful advisor, the voluntary act of the client could also be regarded as some additional effort of the client that will be beneficial to the advisor. The client's contribution to an online feedback platform would be such a possible broader interpretation of our design's voluntary component. Leaving feedback corresponds

30

to an investment of time/effort on the side of the client, while the advisor benefits from a positive rating, at least in an indirect sense.

This is especially interesting since real life transactions in financial/medical advice are not centrally collected by automatized feedback tools (as is the case for online trading platforms like ebay or Amazon). Instead of such an online history that provides 'public' reputation, information about advice quality is essentially limited to own observations and 'private' reputation. However, independent online platforms could gather feedback on a specific type of advice. Clearly, such independent platforms rely more on voluntary contributions of clients than centralized ones as leaving feedback is not just the matter of a mouse click. Our results indicate that some clients are indeed willing to 'pay back' good advice. If clients' potential effort can be channeled into collective feedback, access to 'imperfect public' reputation about the advisor would be possible. As a consequence, market efficiency under real life conditions (finding out on your own about advisors' quality involves transaction costs, quality is multi-dimensional, relatively short time horizons) would benefit.

Generally, online review platforms/systems face a series of challenges.[21] Facilitation of client feedback via a central entity, say, the health insurance system, could help overcome some of these issues. For instance, participation could be institutionalized and promoted by a reduction of the client's insurance premium.

One limitation of our study is that cheating by advisors is modeled to be perfectly detectable. While this can be a realistic feature in some situations, it is not in others. For instance, the low returns from an investment can be due either to the recommendation of an unsuitable financial product or the weak economy. So, accounting for noise by adding a stochastic component which can turn good advice into a bad outcome or bad advice into a good outcome, would extend the scope of our set-up.

---

[21]Naturally, not all clients would be willing to provide feedback on an external site and those who do may not be representative. Often liquidity of feedback seems questionable. See the literature on online feedback platforms, e.g. Lappas et al. (2016), for more details.

Another limitation of our study is that we preclude discriminatory advice and do not allow clients to be recognizable for advisors. In reality, however, advisors are free to give different advice to different clients, whom they typically are able to identify. So, if clients are identifiable and discriminatory advice is allowed, then advisors can retaliate for truthful advice in the past that was not generously rewarded (with, e.g., a bonus) and reward generous (bonus) payments in the past with truthful advice now. In such a set-up, clients would also have incentives to build reputations in order to obtain good advice. Since the opportunities to reciprocate increase, we would expect that the rates of truthful advice will also increase. This aspect remains for future research.

References

Akerlof, G. A., and R. E. Kranton. 2000. "Economics and Identity". *Quarterly Journal of Economics*, **115**, pp 715–753

Anagol, S., S. A. Cole, and S. Sarkar. 2013. "Understanding the Advice of Commissions-Motivated Agents: Evidence from the Indian Life Insurance Market." *Harvard Business School Finance Working Paper No. 12-055. Available at SSRN: http://ssrn.com/abstract=1978876.*

Angelova, V., and T. Regner 2013. "Do voluntary payments to advisors improve the quality of financial advice? An experimental deception game." *Journal of Economic Behavior & Organization* 93: 205–218.

Balafoutas, L., A. Beck, R. Kerschbamer, and M. Sutter. 2013. "What Drives Taxi Drivers? A Field Experiment on Fraud in a Market for Credence Goods." *Review of Economic Studies* 80 (3): 876–891.

Beck A., R. Kerschbamer, J. Qiu, and M. Sutter. 2013. "Shaping beliefs in experimental markets for expert services: Guilt aversion and the impact of promises and money-burning options." *Games and Economic Behavior* 81: 145–164.

Behnk, S., I. Barreda–Tarrazona, and A. Garcia–Gallego. 2014. "The role of ex post transparency in information transmission – An experiment." *Journal of Economic Behavior & Organization* 101: 45–64.

Bénabou, R., and J. Tirole. 2011. Identity, Morals and Taboos: Beliefs as Assets. *Quarterly Journal of Economics*, 126(2), 805–855

Cai, H., and J. T.-Y. Wang. 2006. "Overcommunication in strategic information transmission games." *Games and Economic Behavior* 56(1): 7–36.

Charness, G., and M. Dufwenberg. 2006. "Promises and Partnership." *Econometrica* 74(6): 1579–1601.

Christoffersen, S. E. K., R. Evans, and D. K. Musto. 2013. "What Do Consumers' Fund Flows Maximize? Evidence from Their Brokers' Incentives." *The Journal of Finance* 68: 201–235.

Dana, J., R. Weber, and J.X. Kuang. 2007. "Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness". *Economic Theory*, 33, 67–80.

Danilov, A., T. Biemann, T. Kring, and D. Sliwka. 2013. "The dark side of team incentives: Experimental evidence on advice quality from financial service professionals." *Journal of Economic Behavior & Organization* 93: 266–272.

'Die Welt'. December, 2012. "Falsche Finanzberatung kostet Deutsche Milliarden." *Die Welt, Available at: http://www.welt.de/newsticker/news1/article112238470/Falsche-Finanzberatung-kostet-Deutsche-Milliarden.html*

Dulleck, U., R. Kerschbamer, and M. Sutter. 2011. "The Economics of Credence Goods: An Experiment on the Role of Liability, Verifiability, Reputation, and Competition." *American Economic Review* 101(2): 526–555.

Erat, S., and U. Gneezy 2012. "White Lies." *Management Science*, 58(4):723–733.

Festinger, L. 1957. A Theory of Cognitive Dissonance. *Row Petersen*, Evanston, IL

Fischbacher, U. 2007. "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10(2): 171–178.

Gneezy, U. 2005. "Deception: The Role of Consequences." *American Economic Review* 95(1): 384–394.

Greiner, B. 2004. "An Online Recruitment System for Economic Experiments." In *Forschung und wissenschaftliches Rechnen 2003,* ed. Kurt Kremer, Volker Macho, 79–93. Goettingen: Gesellschaft fuer Wissenschaftliche Datenverarbeitung Bericht 63.

Huck, S., G. K. Lünser, and J.-R. Tyran. 2012. "Competition fosters trust." *Games and Economic Behavior*, 76(1): 195–209.

Huck, S., G. K. Lünser, and J.-R. Tyran. 2016a. "Price competition and reputation in markets for experience goods: an experimental study." *RAND Journal of Economics*, 47(1): 99–117.

Huck, S., G. K. Lünser, F. Spitzer and J.-R. Tyran. 2016b. "Medical Insurance and Free Choice of Physician Shape Patient Overtreatment. A Laboratory Experiment." *Journal of Economic Behavior & Organization*, forthcoming.

Hurkens, S., and N. Kartik. 2009. "Would I Lie to You? On Social Preferences and Lying Aversion." *Experimental Economics* 12(2): 180–192.

Inderst, R., and M. Ottaviani. 2012. "Competition through Commissions and Kickbacks." *American Economic Review*, 102(2): 780–809.

Ismayilov, H., and J. Potters. 2013. "Disclosing advisors' interests neither helps nor hurts." *Journal of Economic Behavior & Organization*, 93: 314–320.

Kerschbamer, R., M. Sutter, and U. Dulleck. 2016. "How social preferences shape incentives in (experimental) markets for credence goods." *The Economic Journal*, in press.

Kreps, D. M., P. Milgrom, J. Roberts, and R. Wilson. 1982. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory*, 27: 245–52.

Lappas, T., G. Sabnis and G. Valkanas. 2016. "The Impact of Fake Reviews on Online Visibility: A Vulnerability Assessment of the Hotel Industry" *Information Systems Research*, in press

Mazar, N., O. Amir and D. Ariely. 2008. "The Dishonesty of Honest People: A Theory of Self-Concept Maintenance". *Journal of Marketing Research*, **XLV**, 633–644.

Mimra, W., A. Rasch and C. Waibel. 2013. "Price Competition and Reputation in Credence Goods Markets: Experimental Evidence". *CER-ETH, Center of Economic Research at ETH Zurich Working Paper No. 13/176.*, available at SSRN: http://ssrn.com/abstract=2240349 or http://dx.doi.org/10.2139/ssrn.2240349

Mullainathan, S., M. Noeth and A. Schoar. 2012. "The Market for Financial Advice: An Audit Study". *NBER Working Paper No. 17929.*

OECD, 2016. "OECD.stat" *http://stats.oecd.org*

Peeters, R., M. Vorsatz, and M. Walzl. 2008. "Rewards in an experimental sender-receiver game." *Economics Letters* 101: 148–150.

Rasch, A., and C. Waibel. 2013. "What Drives Fraud in a Credence Goods Market? Evidence from a Field Study." *CER–ETH – Center of Economic Research at ETH Zurich Working Paper No. 13/179.* Available at SSRN: http://ssrn.com/abstract=2259670 or http://dx.doi.org/10.2139/ssrn.2259670.

Rubin, J., and R. Sheremeta. 2015. "Principal-Agent Settings with Random Shocks." *Management Science* 62(4): 985–999.

Sánches-Pagés, S., and M. Vorsatz. 2007. "An Experimental Study of Truth-Telling in a Sender-Receiver Game." *Games and Economic Behavior* 61(1): 86–112.

Sánches-Pagés, S., and M. Vorsatz. 2009. "Enjoy the Silence: An Experiment on Truth-Telling." *Experimental Economics* 12(2): 220–241.

Schneider, H.S. 2012. "Agency Problems and Reputation in Expert Services: Evidence from Auto Repair." *The Journal of Industrial Economics* 60: 406–433.

Schotter, A. 2003. "Decision Making with Naive Advice." *American Economic Review* 93(2): 196–201.

Vanberg, C. 2015. "Who never tells a lie?" *University of Heidelberg Discussion Paper No. 581*

van de Ven, J., and M. C. Villeval. 2015. "Dishonesty under scrutiny." *Journal of the Economic Science Association* 1(1): 86–99.

6 Appendix

6.1 Data

Table 10: Outcomes by treatment

Treatment	A, F	C, NF	B, F	A, NF	B, NF	C, F	no transaction	interactions
Base	235	0	80	74	18	1	42	450
ID	255	0	119	37	14	2	23	450
Bon	222	0	106	67	23	2	30	450
IDBon	152	0	232	29	7	3	27	450
Comp	190	0	215	20	15	1	9	450
CompBon	115	0	301	11	12	5	6	450
CompID	69	0	365	3	8	0	5	450
CompIDBon	61	0	380	1	6	0	2	450

6.2 Experimental Instructions

Baseline: main text (black)

Identifiability treatments (ID, IDBon, CompID, CompIDBon): additional red text

Bonus treatments (Bon, IDBon, CompBon, CompIDBon): additional green text

Competition treatments (Comp, CompBon, CompID, CompIDBon): additional purple text