Journal of Public Economics 161 (2018) 44–55

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

An offer you *can* refuse: The effect of transparency with endogenous conflict of interest☆

Melis Kartal, James Tremewan*

Department of Economics, Vienna University of Economics and Business, Welthandelsplatz 1, 1020 Vienna, Austria

ARTICLE INFO

Article history:
Received 22 February 2017
Received in revised form 2 April 2018
Accepted 2 April 2018
Available online 24 April 2018

Keywords:
Experimental game theory
Strategic information transmission
Transparency
Conflict of interest
Lobbying

ABSTRACT

We study the effects of transparency on information transmission and decision making theoretically and experimentally. We develop a model in which a decision maker seeks advice from a better-informed adviser whose advice might be swayed by financial incentives. Transparency enables the decision maker to learn whether or not the adviser accepted such an incentive, for example from an "interested" third party. Prior theoretical and experimental research mostly found that transparency is ineffective or harmful to decision makers. Our model predicts that transparency is never harmful and, depending on equilibrium selection, may improve the accuracy of decision makers. In our experiment transparency does indeed improve accuracy, especially if it is mandatory.

1. Introduction

In 2013, the Occupational Safety and Health Administration (OSHA) in the U.S. began a public consultation on setting new limits for working with silica dust, a major health hazard for construction workers that causes serious lung disease. The OSHA created considerable controversy in the Senate when it requested for the first time that those submitting scientific evidence disclose their funding sources. A number of senators protested against the request arguing that revealing this type of information would bias the judgment of the agency. In turn, the head of the OSHA defended the request vigorously, claiming that transparency is indispensable for the information on which the agency bases its decision to meet the highest standard of integrity.[1] How transparency regarding funding sources and financial relationships affects advice and whether it improves the accuracy of decision making in settings such as this, where the expert might be influenced by a third party, is the topic of this paper.

Advice is prevalent in a variety of settings, ranging from regulatory agencies, legislatures, and judiciaries to medical services and financial markets. In such settings, decision makers often face complex decisions with uncertain outcomes, and therefore seek the advice of an expert in order to increase the likelihood of a successful decision. However, information transmission from the expert to the decision maker may be compromised; for example, even if the expert and the decision maker do not have an inherent conflict of interest, a third party, such as a special interest group or an industry, may sway the expert's advice in its favor by offering him a financial reward. Such concern regarding the impartiality of experts funded by third parties was raised in a recent study by the New York Times. The study identified dozens of examples of think tank researchers who helped shape the U.S. government policies in diverse areas such as net-neutrality, military spending, and airport security while being paid by corporations who had stakes in those policies.[2]

On the one hand, transparency is assumed to remedy this kind of situation: it protects the decision maker by revealing whether the expert has a financial incentive that might lead him to give biased

☆ We are grateful to two anonymous referees for helpful comments. We would like to thank Thomas de Haan, Tim Johnson and Dimitri Landa, as well as seminar audiences at GATE-University of Lyon, Heidelberg University, University of Vienna, WZB Berlin Social Science Center, NYU CESS Annual Experimental Political Science Conference, and the 2015 ESA European Meeting.

* Corresponding author.
E-mail addresses: melis.kartal@wu.ac.at (M. Kartal), james.tremewan@wu.ac.at (J. Tremewan).

[1] This was published in Nature on March 4, 2014. In the same issue, an editorial piece argued that regulatory agencies must demand conflict of interest statements for the research that they use.

[2] http://www.nytimes.com/2016/08/09/us/politics/think-tank-scholars-corporate-consultants.html.

https://doi.org/10.1016/j.jpubeco.2018.04.003
0047-2727/

advice. On the other hand, one counter-argument against transparency is that disclosing this type of information results in a bias itself: even if the expert's advice is truthful, the decision maker may dismiss the advice if the expert has a financial tie to an industry or a special interest group.[3] According to the proponents of this idea, the bias against experts funded by interested third parties is harmful, so "the conflict of interest mania" must be cured.[4] An intimately related debate is whether transparency should be *voluntary* or *mandatory*. For example, organizations such as Transparency International advocate mandatory registering of lobbying activity. Most politicians and lobbyists agree that this should be the case, but in countries where registers exist they often remain voluntary.[5,6]

Although it is a widely-held belief that transparency should ameliorate problems related to conflicts of interest, as suggested above it is not immediately clear how it should do so. Moreover, prior theoretical and experimental research has mostly produced bleak results regarding the effects of transparency on decision making. We provide a formal model which illustrates a precise mechanism through which transparency can lead to better decision making. While transparency is never harmful in our model, it does not guarantee strictly better outcomes due to the existence of multiple equilibria. Therefore, we run an experiment to establish whether or not transparency can improve information transmission in practice.

In the model, there are two states of the world (labeled L and R) and two possible policies (labeled l and r). The adviser is an expert who is perfectly informed about the state of the world, whereas the decision maker knows only the prior probability of each state. The adviser recommends a policy to the decision maker, who then makes a policy choice. The payoff of the decision maker is maximized if the chosen policy matches the state of the world. All else being equal, the adviser and the decision maker have no conflict of interest. However, prior to the policy recommendation stage—but after learning the true state—the adviser decides whether or not to accept a side payment: if the adviser accepts the side payment, then he is obliged to recommend policy r.

We consider the following scenarios. In the *mandatory-transparency* condition, the decision maker is informed whether or not the adviser accepted the side payment. In the *non-transparency* condition, the decision of the adviser regarding the payment is not disclosed. We also study a *voluntary-transparency* condition in which transparency is not enforced, and the adviser chooses whether or not to disclose his decision regarding the side payment.

To highlight the basic mechanism through which transparency may improve decision making, we focus on a simple model with stark assumptions, but the mechanism is robust to rich extensions in which these assumptions are relaxed. In particular, it is robust to assuming that the adviser is ex-ante imperfectly informed and obtains higher quality information if he accepts the side payment, and that the adviser is free to choose his recommendation with positive probability even if he accepts the payment.[7]

Our main theoretical result is that transparency (whether it be mandatory or voluntary) never harms decision making and can strictly improve it depending on equilibrium selection — thus, our predictions are not sharp due to equilibrium multiplicity. In order to evaluate whether transparency is indeed beneficial in practice and gain further insights regarding its effect on information transmission and decision making, we designed and ran an experiment on the basis of our model, implementing the three conditions discussed above.

Equilibrium predictions specific to the parameter values that we chose for our experiment are as follows. In the non-transparency condition, the adviser always accepts the side payment and recommends r. As a result, the adviser's recommendation is uninformative. We denote this equilibrium the "corrupt equilibrium." In both transparency conditions, there are two equilibria of interest. The first one is the corrupt equilibrium, in which behavior is the same as in the equilibrium of the non-transparency condition. The second one is what we denote as the "honest equilibrium"; in this equilibrium, the adviser always rejects the payment and gives honest advice, which the decision maker always follows. Sustaining the honest equilibrium requires the type of bias which opponents of transparency argue will be the result of disclosure: if the adviser accepts the side payment, this prompts the decision maker to believe that the adviser is dishonest and choose policy l. This bias is consistent with the "zero tolerance" approach recently adopted by several journals and organizations, which we discuss in more detail in Section 4.1.

Overall, our experiment shows that mandatory transparency clearly improves decision making relative to non-transparency condition. However, the evidence regarding the positive effect of voluntary transparency on decision making is weaker. We find that mandatory transparency improves the accuracy of decisions made in state L, the state in which the adviser has a financial incentive to give a "dishonest" recommendation. While mandatory transparency improves the accuracy of decisions in state L, it has no impact on accuracy in state R. Thus, we conclude that mandatory transparency improves decision making.

The mechanism through which the mandatory-transparency condition improves decision making is consistent with our theory. Many decision makers and advisers view rejecting the side payment as a way to boost the adviser's credibility. When the state is R, many more advisers reject the payment and recommend r in the mandatory-transparency treatment than in the non-transparency treatment. When the state is L, many advisers reject the payment and recommend the correct policy even in the non-transparency treatment—this can be explained by lying-aversion—however, even more advisers do so with mandatory transparency. Thus, advisers' willingness to reject the payment in the mandatory-transparency treatment stems not only from lying-aversion but also from a strategic motive: refusing the side payment is potentially beneficial because if an adviser recommends policy r accepting the side payment, a sizable proportion of decision makers find it suspicious and choose policy l. One caveat is that although the fraction of decision makers mistrusting advisers who accept the payment is nonnegligible, it is also far from a majority. As a result, positive effects of transparency weaken over time: many advisers learn that the negative bias among decision makers against advisers who accept the side payment is not too prevalent and adjust their behavior accordingly.

Our study sheds light on the effects of transparency on adviser and decision maker behavior in an environment where the adviser might be influenced by financial incentives and third party funding. In particular, ours is the first study to combine theory and experiments to show that transparency can help decision makers in such an environment. Transparency is becoming more and more important especially because the share of private enterprise in the funding of research has been rising steeply. According to the National Science Foundation (NSF) in the U.S., the share of industry and government roughly tracked each other until the late 80s. However, industry has since considerably outpaced government in terms of research

[3] See, for example, Stossel (2005), Weber (2009), and Stossel and Stell (2011).

[4] Rago, Joseph. "A Cure for 'Conflict of Interest' Mania." Wall Street Journal, 26 June 2015.

[5] http://www.transparencyinternational.eu/wp-content/uploads/2015/04/Lobbying_web.pdf.

[6] http://www.oecd.org/gov/ethics/Lobbying-Brochure.pdf.

[7] However, we note that mounting evidence suggests that third parties who fund scientists and academic researchers expect the conclusions of their research to be consistent with their interests giving rise to what is known as the "funding bias" in the scientific literature. See Footnote 9 and Section 4.2 for a discussion of highly publicized, large-impact examples.

funding.[8] This has bestowed corporations with an immense influence and ability to shape policy-making as well as public opinion via the research institutions and experts that they fund.[9] In response to this, regulatory agencies, academic journals, NGOs and government agencies have started demanding more transparency. However, this demand for transparency is not without its backlash, as mentioned earlier. Therefore, transparency is not only a very important but also sensitive issue that calls for a thorough theoretical and empirical evaluation.

2. Related literature

2.1. Experimental literature

Prior experimental research mostly found that transparency is either ineffective or has adverse effects on decision makers (Cain et al., 2005, 2011; Koch and Schmidt, 2010; Rode, 2010; Loewenstein et al., 2011, 2012; Ismayilov and Potters, 2013; Behnk et al., 2014).[10,11] Our model differs from the previous experiments—with the exception of Sah and Loewenstein (2014), which we discuss in detail below—in that the conflict of interest between the adviser and the decision maker is endogenous. On the one hand, the adviser and the decision maker have no conflict of interest if the adviser rejects the side payment. Hence, transparency allows the adviser to send a message to the decision maker that his advice is honest by rejecting the side payment. On the other hand, accepting the payment per se is not inherently dishonest as explained below. Whether or not the adviser accepts the side payment and how the decision maker interprets the advice of an adviser who has accepted the payment are determined in equilibrium.

Sah and Loewenstein (2014) is the only other experimental study which endogenizes the conflict of interest in information transmission, albeit using a different approach than ours. In their design, the adviser decides whether or not to accept a conflict of interest, and transparency enables the decision maker to learn whether the adviser has accepted being conflicted, in which case he always has an incentive to lie to the decision maker. Our design differs from Sah and Loewenstein (2014) in that accepting a side payment is not equivalent to accepting a conflict of interest; as a stark example, an adviser who accepts the payment *only* in state *R* is never conflicted in our model. Therefore, the implications of disclosing that the adviser accepted a side payment are both theoretically and experimentally more subtle than those of disclosing that there is an outright conflict of interest. Our approach is more reflective of the situations that we are interested in where, for example, researchers funded by a pharmaceutical company will not necessarily misreport findings. In line with our findings, Sah and Loewenstein (2014) documents that transparency increases the accuracy of decision making and find no evidence of moral licensing. Unlike our study, they find no difference in the improvement in the accuracy of decision-making between voluntary and mandatory transparency — this is likely due to the difference in the designs, which we discussed above.

Our model and experimental results also relate to studies on lying aversion. Previous experiments have shown that people may find lying morally costly and avoid it (see, for example Gibson et al., 2013; Gneezy, 2005; Hurkens and Kartik, 2009; Sánchez-Pagés and Vorsatz, 2007). Our results are consistent with previous findings in that many advisers choose to tell the truth even if there is a material incentive to lie.

2.2. Theoretical literature

There is an extensive literature on strategic information transmission from a better-informed sender to a receiver, dating back to the seminal works by Crawford and Sobel (1982) and Sobel (1985). In this literature, the receiver and the sender have payoffs that are misaligned to a certain degree. This misalignment results in a bias in the sender's communication; the sender's equilibrium message to the receiver is typically noisy and can even be uninformative depending on the precise structure of payoffs.

In most of the literature on strategic information transmission, the extent to which sender and receiver payoffs are (mis)aligned is exogenous. Our model endogenizes it since the adviser may choose to decline the side payment. The extent to which the adviser and the decision maker payoffs are aligned is also endogenous in Durbin and Iyer (2009), and Inderst and Ottaviani (2012a,b), but this is due to strategic third parties. There is a single third party in Durbin and Iyer (2009) and Inderst and Ottaviani (2012b), whereas there are multiple third parties in Inderst and Ottaviani (2012a). In these models, third parties set side payments for the adviser (sales commissions, bribes, etc.) in order to maximize their return. Inderst and Ottaviani (2012a,b) also analyze the effects of transparency in their setting. Inderst and Ottaviani (2012a) show that transparency can have adverse effects on the decision maker, unlike in our model. In Inderst and Ottaviani (2012a), third parties produce horizontally-differentiated products and compete for consumers using advisers who are incentivized by sales commissions. We abstract from the competition of third parties and focus on the effects of transparency with only one possible source of bias: this can be thought of as a single or disproportionately powerful third party (e.g., oil and energy lobby, tobacco industry, sugar industry, and gun rights lobby) that can affect policy and sway public opinion through political lobbying, and funding research institutes and experts. Inderst and Ottaviani (2012b) analyze a model of financial advice and show that transparency can be beneficial in their setting provided that the decision maker is naive, and disclosing the commission payments to the adviser turns the naive decision maker into a wary one. To our knowledge, this is the only other model in which transparency can be beneficial, and our study complements it by showing that transparency can be beneficial even with a rational decision maker and corroborating theoretical results with experimental evidence.

Also related is the model by Potters and Van Winden (1992, 2000) in which there are two players, a policy-maker and a better-informed interest group (i.e., a lobby) that can send the policy-maker a costly message. Their model differs from the cheap talk literature since sending a message entails a cost to the lobbyist — this relates to lobbying costs. Our model can be thought of as a lobbying model (absent lobbying costs) in which financial incentives can influence the message of the lobbyist to the policy-maker.

[8] http://www.nsf.gov/statistics/seind14/index.cfm/chapter-4/c4s1.htm.

[9] For example, the energy industry, the tobacco industry, and the sugar industry have used experts extensively in order to suppress information regarding the hazards of their products. One highly publicized case involves ExxonMobil, the world's largest oil company. ExxonMobil made the headlines recently due to its funding of climate change denial. It was reported that the company knew as early as 1981 of climate change and despite this, spent a substantial amount of money over the next three decades to promote its denial (The Guardian, July 8, 2015). On a related note, only 47% of US voters believe that climate change is caused by human activities, compared with 97% of climate scientists (Yale Project on Climate Change Communication, 2013 — http://environment.yale.edu/climate-communication-OFF/files/ClimateNote_Consensus_Gap_May2013_FINAL5.pdf).

[10] Loewenstein et al. (2011) put forth two explanations for the adverse effects of transparency that they observe. Firstly, transparency may result in "moral licensing": when conflicts of interest are disclosed to decision makers, advisers may feel less obliged to act morally and give accurate advice than when the decision makers remain in the dark — we find no evidence of this effect in our experiment. Secondly, decision makers fail to sufficiently discount recommendations by advisers whose biases are disclosed. Via these two channels, decision makers obtain worse outcomes with transparency.

[11] In Behnk et al. (2014), transparency reduces the frequency of deceptive messages in only one of their six experimental conditions (whether decision maker accuracy benefits from transparency was not the focus of the study).

3. The model

We develop a model with two players: an adviser (A) and a decision maker (D). There are two possible states of the world. Nature draws the state $S \in \{L, R\}$ such that

$$S = \begin{cases} L, & \text{with probability} \quad p \in (0,1) \\ R, & \text{otherwise.} \end{cases}$$

The prior probability p is common knowledge. The adviser (A) learns the true state of the world, whereas the decision maker (D) knows only the prior p. After learning the state, A recommends a policy $s \in \{l, r\}$ to D, who then chooses a policy. D prefers the policy to match the state of the world. After learning the state and before making a recommendation, A decides whether or not to accept a "side payment." If A accepts the side payment, then he *must* recommend policy r, whereas if A rejects the payment, then he decides whether to recommend l or r. Note that an adviser who accepts the payment is not necessarily conflicted: A may choose to accept the side payment only in state R and remain honest.

The side payment can be interpreted as funding, gifts, or secondary employment and consulting fees by a third party, such as a special interest group or an industry which strictly prefers policy r regardless of the state of the world, or as the possibility to invest in the stock or equity of a corporation that will benefit from r. To give real-life examples (see also the discussion in Section 4.2), a nutrition scientist who investigates the effect of a high-sugar diet on coronary heart disease decides whether or not to accept funding from the sugar industry, or an expert whose testimony is sought regarding the practice of an industry decides whether or not to have (or stop having) a financial tie to that industry.[12]

In Section 4.2, we discuss our modeling assumptions and their relevance in light of the evidence regarding industry-funded research. In that section, we also show that our main results are robust to relaxing these assumptions. In particular, they are robust to assuming that the adviser initially obtains an imperfect signal regarding the true state and obtains higher quality information only if he accepts the side payment, and that the adviser is free to choose his recommendation with some positive probability even if he accepts the payment. We consider the following scenarios in our analysis.

(i) In the "mandatory-transparency condition," D learns whether or not A accepted the side payment before choosing the policy.
(ii) In the "voluntary-transparency condition," A decides whether or not to disclose his decision regarding the side payment (after learning the state).
(iii) In the "non-transparency condition," D has no information about the decision of A regarding the payment.

While an adviser cannot actually be forbidden to make a claim about his decision regarding side payments, we argue that unless a transparency policy is set in place and institutionalized, making a claim that one has not accepted a side payment cannot be easy to verify. As discussed in Chen et al. (2017), data on disclosed activity rarely exist in the absence of extensive disclosure laws. For example, before the "sunshine laws" were widely enacted in the US, it was virtually impossible to determine whether or not a physician or researcher accepted payments or gifts from the industry, and before peer-reviewed journals started implementing transparency policies, few articles contained financial disclosures. Thus, the scarcity of data on financial relationships before a transparency policy is implemented implies that it is unlikely that either an adviser admits accepting side payments without mandatory transparency or that a claim that a side payment is not accepted is verifiable.

3.1. Payoffs

D obtains a payoff of $\bar{\pi}$ if the chosen policy matches the state of the world, and a payoff of $\underline{\pi}$ otherwise, where $\underline{\pi} < \bar{\pi}$. A receives

(i) $\alpha > 0$ if D chooses the policy that A recommends, and
(ii) $\gamma \geq 0$ if A recommends the "better" policy for D, and
(iii) $\beta_S > 0$ if A accepts the payment in state $S \in \{L, R\}$.

We assume that $\alpha > 0$ because being followed benefits adviser reputation and enables future business prospects. If for example the adviser is a lobbyist, then α represents an increase in earnings, as a lobbyist with a track record of influencing policymakers will attract more clients. If the adviser is a scientist, then α represents impact and recognition in the scientist's field, research grants, publications, professorial positions and journal editorships.

Our modeling approach reflects situations where there may be a significant lag between the advice and the resolution of uncertainty regarding the true state as in our silica dust regulation example. The resolution of uncertainty will take (or has taken) decades in numerous, important circumstances — consider the relationship between tobacco use and cancer, the relationship between high-sugar diet and coronary heart disease, the climate change debate, the current debate on asylum and immigration policies, the debate on genetically modified organisms, the debate on affordable health care policies, and so on. This means that if an expert chooses to provide advice that goes against his information, he cannot easily be held accountable.[13] Hence, α represents the payoff from being consulted and followed, irrespective of the consequences, as it may take a long time to evaluate the accuracy of advice — in some cases, even longer than the expert's lifetime.

Unlike α, the parameter γ captures the intangible consequences of A's recommendation. While assuming that γ is strictly positive is not necessary for our results, we allow for this possibility due to reasons such as lying aversion or the cost of manipulating research findings.[14] We also allow for state-dependent side payments. Again, this is not necessary for our results; however if for example $\gamma > 0$, then it may be reasonable to allow for the case in which $\beta_L > \beta_R$ because L is the state in which an adviser who accepts the side payment has to give a dishonest recommendation and lose γ.[15]

[12] One can also envision a model in which the adviser makes a decision about the side payment before conducting a full-fledged research to find out the true state. Our main results are robust to assuming that the adviser has an ex-ante imperfect signal and must decide about the payment before obtaining the complete results of his research as long as the accuracy of his signal is higher than the accuracy of the prior. In many cases, it is reasonable to assume that the adviser has some ex-ante information which is more precise than what the public knows (embodied in the prior) by virtue of being an expert. Going by the examples cited above in the main text, a nutrition scientist likely knows more about the effect of high-sugar diet, and an expert who may testify in court regarding the conduct of an industry is more knowledgeable about the industry than the public, even before making a decision regarding the side payment.

[13] On a related note, a meta-study of (anonymous) scientist surveys by Fanelli (2009) found that on average, over 14% of respondents reported to have observed fabrication, falsification and modification of data by colleagues, and up to 72% have observed other questionable practices. However, Fanelli (2009) notes that fabrication and falsification are very hard to detect in the data, and are rarely reported by whistle blowers.

[14] Alternatively, γ could represent having a concern for the outcome in the sense that A receives γ if D chooses the correct policy. In that case, all of our results go through with a minor tweak.

[15] Note that β_L and β_R can differ if, for example, a special interest group or corporation has its in-house researchers and knows the state, but wants to employ an outside expert to use the expert's credibility; this is indeed common practice for the tobacco industry, pharmaceutical companies, and the energy industry (see in particular the case of ExxonMobil in Footnote 9).

4. Equilibrium analysis

Presenting the analysis of all three conditions in full generality requires a considerable increase in length and complexity of exposition with little corresponding increase in insight. Therefore, we focus on the analysis of mandatory-transparency and non-transparency conditions in the main text; Online Appendix A presents the main results for the voluntary-transparency condition in detail, and Section 6.4 presents a brief overview of the key results. Hereafter, "transparency" refers to mandatory transparency unless otherwise stated.

Let $a \in \{r_A, r_R, l\}$ denote A's action, where $a = r_A$ if A accepts the side payment and recommends r, $a = r_R$ if A rejects the payment and recommends r, and $a = l$ if A recommends l (as explained above, A *must* reject the side payment in order to be able to recommend l). Next, let m denote A's "message" to D. In the transparency condition, D fully observes A's action. So, A's message is equivalent to his action, and $m = a$. In the non-transparency condition, A's message is equivalent to his policy recommendation. If $a \in \{r_A, r_R\}$, then $m = r$. If $a = l$, then $m = l$. Hence, the message space is $\{r_A, r_R, l\}$ ($\{r, l\}$) in the transparency (non-transparency) condition.

If $\gamma > 0$, then accepting the side payment results in dishonest advice and loss of γ for A *only* in state L. Therefore, we assume that $\beta_L > \gamma$ holds as otherwise giving honest advice in state L is sufficiently important to the adviser that the "agency problem" is virtually nonexistent under both transparency and non-transparency conditions; that is, full information transmission is an equilibrium outcome in both conditions.

Assumption 1. (Existence of an agency problem). $\beta_L > \gamma$.

We also maintain the assumption that $\alpha \geq \beta_R$. If $\alpha < \beta_R$ and $\beta_L > \gamma$, then side payments are too high, and the adviser cannot be expected to behave in the best interest of the decision maker in either condition. Therefore, we impose Assumption 2 in the subsequent analysis.

Assumption 2. (Importance of adviser reputation). $\alpha \geq \beta_R$.

Assumption 2 reflects the case where reputation and being recognized as a reliable source of judgment are sufficiently important for the adviser. Indeed, such reputation is integral to the career of advisers — for example, we expect α to be sufficiently high for most scientists because a scientist needs to, among other things, continually publish in peer-reviewed journals for professional recognition (i.e., a scientific publication and professional recognition are analogous to being followed and receiving α, respectively). In a similar vein, a lobbyist who has a track record of recent successful campaigning and influencing policy makers is more likely to reach new clients. In what follows, we describe the results derived from the analysis of the perfect Bayesian Equilibria (PBE) of the transparency and non-transparency conditions under Assumptions 1 and 2. A formal description of the PBE is relegated to Online Appendix B.1.

4.1. Results

Since our model gives rise to various types of equilibria depending on the precise parameters, in this section we only provide a general set of theoretical results with the aim of providing a welfare comparison of the transparency and non-transparency conditions. We explain in detail PBE specific to our experimental sessions in Section 5.1, and prove our equilibrium statements in Online Appendix B.2.

The "normative" measure of interest in our comparative analysis is the accuracy of decision making. We choose to focus on the accuracy of decision making rather than the efficiency of the aggregate expected payoff because in those cases in which we are interested the payoff to the adviser is insignificant relative to the social costs of implementing the wrong policy. We first provide a result for the non-transparency game: the presence of an adviser does not improve the ex ante accuracy of decision making relative to the decision maker acting based only on her prior.

Proposition 1. *Under Assumption 1, the ex-ante accuracy of decision making in the non-transparency condition is equal to* $\max\{p, (1-p)\}$ *in every PBE.*

In the transparency game, in addition to A's advice, D observes whether or not A accepted the side payment. Under Assumption 2, the accuracy of decision making is at least $\max\{p, (1-p)\}$ in every PBE, and there always exists a PBE in which D makes the correct decision with probability one in both states. However, there is equilibrium multiplicity as there also exists an equilibrium with an accuracy of $\max\{p, (1-p)\}$.

Proposition 2. *Under Assumption 2, in the transparency condition (i) the ex-ante accuracy of decision making is weakly greater than* $\max\{p, (1-p)\}$ *in every PBE, (ii) there always exists a PBE in which the accuracy is equal to 1, and (iii) there always exists a PBE in which the accuracy is equal to* $\max\{p, (1-p)\}$.

Propositions 1 and 2 are robust to the Intuitive Criterion (Cho and Kreps, 1987), with the exception of Proposition 2 part (iii), as this refinement sometimes uniquely selects the fully informative equilibrium under transparency (we provide an example at the end of the proof of Proposition 2 part (iii) in Online Appendix B.1). Combining these findings we obtain the main result of our model: transparency never hurts and, depending on equilibrium selection, can strictly improve the accuracy of decision making.[16]

Proposition 3. *Under Assumptions 1 and 2, the ex-ante accuracy of decision making in the transparency condition is at least as high as that in the non-transparency condition. Moreover, under transparency there always exists a PBE with an accuracy of 1, exceeding the accuracy of any PBE under non-transparency.*

The intuition for our main result is as follows. In every PBE in which transparency strictly improves decision making, there is a severe bias against an adviser who chooses r_A; that is, the decision maker is committed to punishing an adviser who accepts the payment and recommends r by not following his recommendation. This bias is effective and results in full information transmission because by Assumption 2 the adviser sufficiently cares about his reputation and his influence on the decision maker. The non-transparency condition does not allow for this type of bias against an adviser who chooses r_A, and thus, advice does not improve the accuracy of decision making — indeed, transparency also does not improve accuracy in its equilibria without such a bias. The bias against the acceptance of the side payment is consistent with the "zero-tolerance" approach recently adopted in a number of journals and organizations regarding memberships, peer-reviewed publications, editorials

[16] Proposition 3 is robust to the Intuitive Criterion, see Online Appendix B.1 for its definition adapted to our model. There, we also define D1 (Cho and Kreps, 1987), show that the first statement of Proposition 3 is always robust to D1, and characterize the conditions under which the second statement is robust.

and editorships. American Thoracic Society journals, the BMJ (formerly the British Medical Journal), BMJ Open, Thorax and Heart, PLoS Medicine, PLoS One, PLoS Biology, and the Journal of Health Psychology refuse to publish research partly or fully funded by the tobacco industry. The American Society of Gene Therapy calls for the prohibition of financial relationships between its members and interested third parties. The New England Journal of Medicine (NEJM) requires that authors of editorials do not have financial ties to companies that make products related to the issues they discuss "since editorialists do not provide data, but instead selectively review the literature and offer their judgments." Moreover, none of the NEJM editors can have any financial relationship with any biomedical company. As written in an editorial by Marcia Angell (then editor-in-chief of NEJM) in May 2000, NEJM does "not believe disclosure is enough to deal with the problem of possible bias."[17]

4.2. Discussion

We close this section with a discussion of richer variations of our model. It may be argued that in certain fields of expertise, the adviser may use third-party funding in order to acquire more accurate information. Allowing for this possibility cannot affect any of our results unless we relax the assumption that the adviser who accepts the side payment must recommend r. However, our results are robust to jointly assuming that accepting the side payment increases the accuracy of the adviser's information, and an adviser who accepts the payment can choose between recommending l or r with a positive—but not too high—probability.[18]

How plausible is it to assume that an adviser who accepts a payment from a third party will have to recommend the policy that the third party favors with a high enough probability? We believe that it is a reasonable assumption in many circumstances given the body of evidence accumulated so far regarding what is known as the "funding bias" in the literature. In a very recent study, Kearns et al. (2016) documented that sugar industry executives funded prominent nutrition scientists in 1960s in return for review articles that would falsify studies which pointed out that a high-sugar diet posed a major risk for coronary heart disease (CHD). Moreover, the sugar industry successfully cast doubt about the hazards of sugar for decades while promoting fat as the main dietary risk for CHD. These findings led Kearns et al. (2016) to argue that "policy making committees should consider giving less weight to food industry-funded studies." Another highly-publicized case involves tobacco companies, which have been repeatedly sued both for fraud—hiding from the public what they knew about their product—and in order to recover health costs associated with smoking. The tobacco industry employed dozens of experts in order to testify on their behalf. Kenneth Ludmerer, a distinguished professor of history and medicine at Washington University in St. Louis testified as an expert on medical history on behalf of the tobacco industry over a period of 15 years. From the testimony of Ludmerer in 2002:

Question: Doctor, is it your opinion that cigarette smoking contributes to the development of lung cancer in human beings?

Answer: I have no opinion on that.

Ludmerer's testimony implies that tobacco companies cannot be held liable for any wrongdoing since he, the expert, has no opinion on whether cigarette smoking contributes to the development of lung cancer — as recently as in 2002. Ludmerer was paid more than $550,000 by the tobacco industry (Delafontaine, 2015). It seems unlikely that the industry would pay this amount to an expert who could give an affirmative answer to the question above. On the contrary, Robert Proctor and Louis Kyriakoudes, two (of only three) experts who testified against tobacco companies were subject to harassment by the industry (Delafontaine, 2015).[19]

5. Experimental design

As discussed in Section 4.1, our equilibrium results are typically not sharp due to equilibrium multiplicity. To gain further insights regarding the behavioral effects of transparency, we designed and ran an experiment which implemented our model. Our aim is to empirically answer the following research questions:

1) Does transparency help decision makers?
2) Do advisers behave differently if decision makers observe their decision regarding the side payment? In particular, are advisers less likely to accept the payment in order to enhance their credibility?
3) Do decision makers take into account or ignore advisers' decision regarding the side payment when they choose the policy?
4) Does the accuracy of decisions depend on whether transparency is mandatory or voluntary?

5.1. Experimental parameters

The parameter values we used in the experiment are as follows: $p = 0.4, \alpha = 6, \beta_L = 5, \beta_R = 2, \gamma = 1, \bar{\pi} = 10, \underline{\pi} = 5$. In our equilibrium analysis, we focus on equilibria that satisfy the Intuitive Criterion (IC).[20] Given these parameters, there exists a unique PBE in the non-transparency condition, and we denote it as the "corrupt equilibrium." In the corrupt equilibrium, A always chooses r_A, and D always follows A's advice and chooses r. Intuitively, the corrupt equilibrium arises because A's message is uninformative, and thus D relies on her prior to choose a policy. It is in the best interest of D to choose r as $p = 0.4$. As a result, an adviser who chooses r_A attains the highest possible payoff in both states and has no incentive to deviate.

There exist three PBEs in the transparency condition given the parameter values above. The first one is the corrupt equilibrium in which A always chooses r_A and D follows A's advice as in the corrupt equilibrium of the non-transparency condition. The second is what we denote the "honest equilibrium." The honest equilibrium is fully informative: A always rejects the payment and provides honest advice, and D follows A's advice. The third is a mixed-strategy equilibrium; we do not dwell on a detailed explanation of this equilibrium as we find it unlikely to arise in practice.[21] Table 1 presents the equilibrium predictions for accuracy by state and equilibrium type.

The mechanism of the corrupt equilibrium is identical to that in the non-transparency condition. The intuition for the honest

[17] On a related note, Kesselheim et al. (2012) show that industry sponsorship of a clinical trial negatively influences the perception of physicians regarding the methodological quality of the trial and reduces their willingness to believe and act on trial findings, consistent with the bias against an adviser who chooses r_A we described above.

[18] Assume, as an example, that an adviser who accepts the side payment is perfectly informed about the state, whereas an adviser who does not accept it gets a signal $\sigma \in \{L, R\}$ about the state such that $\Pr(\sigma = S|S) = 0.8$. Assume also that $p = 0.4$ as in our experimental design. If the adviser who accepts the side payment is free to choose the policy recommendation with a probability that is lower than 0.39, then our main result goes through in the sense that transparency does not undermine decision making, and it still generates the most informative equilibrium.

[19] http://www.thenation.com/article/big-tobacco-and-historians/.

[20] Formal equilibrium statements and their proofs can be found in Online Appendix B.2.

[21] In this equilibrium, A always chooses r_A in state L and randomizes between r_R and r_A in state R, and D randomizes between l and r if A chooses r_A and chooses r if A chooses r_R. In line with our expectations, we found no evidence of equilibrium mixed-strategy play: only three out of 72 advisers played in proportions close to the equilibrium prediction.

Table 1
Equilibrium predictions for decision maker accuracy by state and equilibrium.

Equilibrium type	State L	State R
Corrupt equilibrium	0	1
Honest equilibrium	1	1
Mixed-strategy equilibrium	0.33	0.67

Table 2
Proportion of correct decisions by state and treatment.

	Transparency	Non-transparency	Mann-Whitney p–value
state L	0.461	0.297	0.008
state R	0.756	0.793	0.337

Note: Mann-Whitney tests using matching-group averages as observations ($n = 12$).

equilibrium is consistent with our discussion in Section 4.1. In the honest equilibrium, A is deterred from accepting the payment in either state because D is biased against an adviser who chooses r_A; that is, D is committed to punishing an adviser who accepts the payment by not following his recommendation. This bias results in full information transmission as α is sufficiently high; in equilibrium, A always rejects the side payment and gives honest advice. These equilibria satisfy not only the Intuitive Criterion (IC) but also D1 (Cho and Kreps, 1987) as shown in Online Appendix B.2.[22]

5.2. Experimental protocol and hypotheses

We ran three treatments implementing the (mandatory-) transparency, non-transparency, and voluntary-transparency conditions. The experiment consisted of 40 rounds. We used a between-subject design with stranger matching. At the beginning of the experiment, each subject was randomly assigned to be a receiver (i.e., decision maker) or a sender (i.e., an adviser). Subjects remained in the same role for the first 20 rounds. After 20 rounds were over, the roles were switched and subjects remained in their new role until the end of the experiment. Subjects were not informed at the beginning of the experiment that roles would be switched after the first 20 rounds.[23]

Before the start of the experiment and the assignment of roles, subjects went through a tutorial and answered control questions. Decision making in each round of the experiment was as described in the theory section. We used the strategy method with decision makers in order to obtain more observations at each information set and analyze better whether subjects use equilibrium strategies in this role.

Instructions involved neutral language and stated that the sender could choose to accept an extra payment, in which case the sender had to recommend r.[24] The experiment was conducted at the experimental laboratory of the Vienna Center for Experimental Economics (VCEE) at the University of Vienna. Subjects were recruited from the general student population in Vienna via ORSEE (Greiner, 2015). We ran three sessions for each treatment. In each session, there were 24 subjects, each assigned to one of two matching groups, giving us six independent observations for each treatment. All sessions were conducted using a computer program written in z-Tree (Fischbacher, 2007). Sessions lasted around 75 min and the average payoff per subject was approximately 16 euros (including a fee for completing the questionnaire at the end of the experiment).

We now present the two experimental hypotheses that we derive from our theoretical results given the parameters used in the design. Hypothesis 1 parts (a) and (b) concern the accuracy of decision making by state and takes into account all possible, plausible equilibria given the realized state.

Hypothesis 1. *(a) Transparency increases decision maker accuracy in state L. (b) Transparency has no effect on decision maker accuracy in state R.*

Part (a) is due to the following. L is the state in which A accepts the side payment and gives dishonest advice in the corrupt equilibrium. Thus, corrupt equilibrium strategies result in the wrong decision in state L. The corrupt equilibrium is the unique prediction in the non-transparency treatment, whereas we expect a fraction of subjects to play according to the honest equilibrium in the transparency treatment, resulting in correct decisions in state L. Part (b) follows because if the state is R, then D always makes the correct decision in both the corrupt equilibrium and the honest equilibrium. While we do not expect mixed-equilibrium strategies to be empirically relevant (this is indeed confirmed in our data), assuming that a positive fraction of subjects adopt such strategies when deriving our hypotheses would make no difference to the conclusions we draw from our experiment.[25] Finally, we have the following hypothesis because we predict a fraction of subjects in the adviser role to behave in accordance with the honest equilibrium.

Hypothesis 2. *Honest equilibrium behavior is more prevalent under transparency in both L and R — i.e., transparency increases the fraction of advisers who reject the payment and recommend the correct policy in both states.*

6. Results

In this section, we address each of our research questions in turn. When testing our hypotheses we use one-sided tests when we have a clear directional hypothesis, and two-sided tests otherwise.

6.1. Does transparency help decision makers?

In this section, we analyze the accuracy of decisions by treatment. We start with state L. Hypothesis 1 (a) states that transparency increases decision maker accuracy in state L, the state in which an adviser who accepts the side payment will give dishonest advice. Table 2 displays the proportion of correct choices by state and treatment: across all the rounds, 46.1% of decisions in state L are correct in the transparency treatment, whereas only 29.7% are correct in the non-transparency treatment. This difference of 16.4% is statistically significant according to a one-sided Mann-Whitney

[22] For the transparency treatments to have the best shot at improving experimental decision making, we chose parameter values such that the honest equilibrium satisfies not only IC but also D1, which is a more stringent condition than IC. We ran pilot sessions in which all the parameter values were the same except that $\beta_R = \beta_L = 4$. In this case, the honest equilibrium is not robust to D1, and in the experiment, transparency did not have a statistically significant impact on decision-making, and it had either zero or little impact on other dimensions. Therefore, we moved forward with treatments in which the honest equilibrium is robust to both IC and D1. See the more comprehensive discussion concerning the previous and current treatments in the Conclusion.

[23] We applied role-switching because it may facilitate learning.

[24] Instructions can be found in Online Appendix C.

[25] More precisely, assuming that a positive fraction of subjects adopt mixed-equilibrium strategies makes no difference in Hypothesis 1 (a) because such strategies also increase decision making accuracy relative to the corrupt equilibrium in state L as indicated in Table 3. As for state R, such strategies decrease decision making accuracy relative to other equilibria, and the alternative to Hypothesis 1 (b) as follows: "Transparency reduces decision maker accuracy in state R." However, the statistical conclusion remains the same as we cannot reject the null with either formulation.

Table 3
Percentage of adviser behavior consistent with honest equilibrium.

	Transparency	Non-transparency	Mann-Whitney p–value
State L			
Overall	0.303	0.224	0.075
Before role-switch	0.313	0.217	0.027
After role-switch	0.294	0.230	0.315
State R			
Overall	0.151	0.027	0.039
Before role-switch	0.195	0.028	0.017
After role-switch	0.104	0.026	0.129

Note: One-sided Mann-Whitney tests using matching-group averages as observations ($n = 12$).

test ($p = 0.008$). We conclude that transparency improves the accuracy of decision makers in state L in line with Hypothesis 1 (a). Accuracy in state L in the non-transparency treatment is larger than the theoretical prediction of zero mainly due to the fact that even in the non–transparency treatment there are many advisers who reject the side payment and recommend the correct policy in state L (this is likely due to lying-aversion and/or social preferences). Most decision makers follow these advisers and make the correct decision in state L.

Next, we study the accuracy of decisions in state R, the state in which even an adviser who accepts the side payment will give honest advice. Table 2 shows that 75.6% of decision makers choose policy r if the state is R in the transparency treatment. The percentage in the non-transparency treatment is 79.3%. The difference between the two treatments is not significant ($p = 0.337$ according to a two-sided Mann-Whitney test), consistent with Hypothesis 1 (b).

Overall, our answer to Question 1 in Section 5 is affirmative as transparency makes decision makers better off in state L and has no effect in state R. Finally, we note that decision maker accuracy remains more or less constant over time in both treatments.

6.2. Adviser behavior

Hypothesis 2 predicts that transparency increases the fraction of advisers who behave in accordance with the honest equilibrium (i.e., choose l in state L and choose r_R in state R). The relevant data is summarized in Table 3.

We start by analyzing A's behavior in state L. Table 3 shows that the fraction of advisers who reject the payment and recommend l in state L is 30.3% in the transparency treatment and 22.3% in the non-transparency treatment.[26] The difference across the two treatments is weakly significant according to a one-sided Mann-Whitney test as Table 3 shows ($p = 0.075$). However, if we consider the data before and after role-switching separately, the treatment difference is statistically significant at 5% level in the first 20 rounds ($p = 0.027$) and insignificant in the second. We observe a similar trend in state R. We will discuss possible reasons for this trend in adviser behavior shortly.

In the non-transparency treatment, 2.7% of advisers choose r_R in state R compared to 15.1% with transparency. The difference across the two treatments is significant at the 5% level ($p = 0.036$). As in state L , the effect of transparency is stronger in the first half of the experiment; on average, 19.5% of advisers reject the payment and recommend r in state R with transparency, whereas only 2.8% do so without transparency ($p = 0.017$).

From these findings, we conclude that transparency increases the percentage of advisers who reject the payment and recommend the correct policy in both states, in line with Hypothesis 2. Thus, our answer to Question 2 is also affirmative: transparency does affect advisers' behavior and makes them more likely to follow the honest equilibrium strategy.

One caveat is that the effect of transparency on adviser behavior weakens over time, as mentioned above. Figs. 1 and 2 display the evolution of adviser behavior consistent with the honest equilibrium in the transparency treatment in states L and R. The figure reveals a downward trend in both states; the trend is especially prominent in state R.

What are the plausible explanations for this observation? As we discuss in the next section in more detail, a substantial fraction of decision makers in the transparency treatment choose to follow the recommendation of the adviser even if they learn that the adviser accepted the side payment. This behavior is likely to reduce over time the proportion of advisers who behave in accordance with the honest equilibrium. Indeed, our regression analysis in the next section shows that advisers' behavior is shaped by their previous experience with decision makers and that advisers who previously faced limited or no bias for accepting the side payment are more likely to accept the payment than those who have often been punished for it.

We also note that our one-time role switching rule may have had a negative influence on the decision makers of the first half of the experiment, who became advisers in the second half. For instance, decision makers in the first half of the transparency treatment were recommended the correct policy less than 1/3 of the time in state L, as Table 3 shows. This type of behavior may have resulted in resentment among decision makers and facilitated the pervasiveness of the corrupt equilibrium behavior in the second half when the decision makers of the first half became advisers to their previous advisers.

6.3. Decision maker behavior

We start with the non-transparency treatment. The corrupt equilibrium is the unique equilibrium in this treatment. To be more precise, the equilibrium is unique on the equilibrium path but there are different off-the-equilibrium path beliefs and decision maker strategies that support the same equilibrium that cannot be eliminated by IC or D1. Out of four possible pure strategies, two are consistent with the corrupt equilibrium: (i) always choose r; and (ii) choose r if $m = r$ and l if $m = l$. Adopted by 63.8% of decision makers, the strategy "choose r if $m = r$ and l if $m = l''$ is by far the most popular strategy. Fig. 3 shows the distribution of decision maker strategies in the non-transparency treatment. In total, 82.6% of the decision makers use strategies that are consistent with the corrupt equilibrium.

We next find the best response of the decision maker in the non-transparency treatment given the empirical distribution of the adviser recommendations. This is equivalent to finding the empirical posterior probability of state R given that (i) the adviser recommends r; and (ii) the adviser recommends l. We find that the optimal decision maker strategy is "choose r if $m = r$ and l if $m = l$," which is indeed the most frequently observed strategy in the data.

Fig. 4 shows the distribution of decision maker strategies in the transparency treatment. Of the eight possible decision maker strategies, "choose r if $m \in \{r_A, r_R\}$ and l if $m = l''$ is the most popular one — about 59% of decision makers use this strategy. This strategy is consistent with the corrupt equilibrium. More generally, there are four decision maker strategies that are consistent with the corrupt

[26] Note that in both treatments a small percentage of advisers reject the payment and recommend the wrong policy nevertheless. This occurs in both state L (2.3% with transparency; 1.4% with non-transparency) and state R (3.3% with transparency; 5.7% with non-transparency).

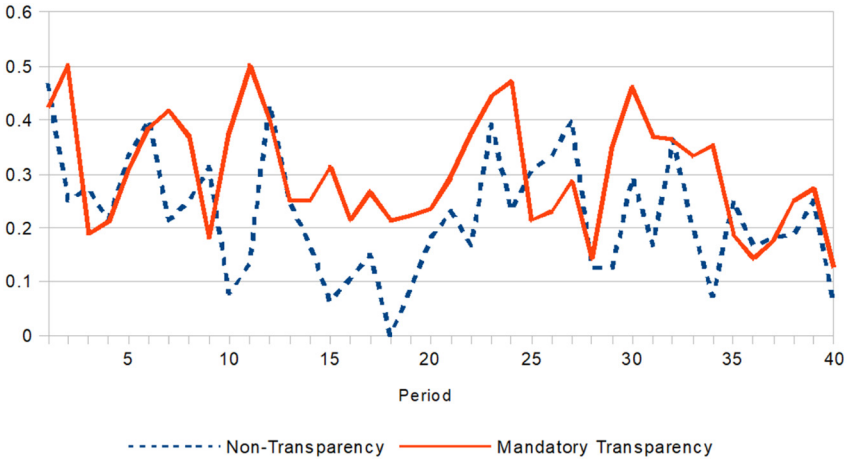

Fig. 1. Percentage of advisers who recommend l in state *L*.

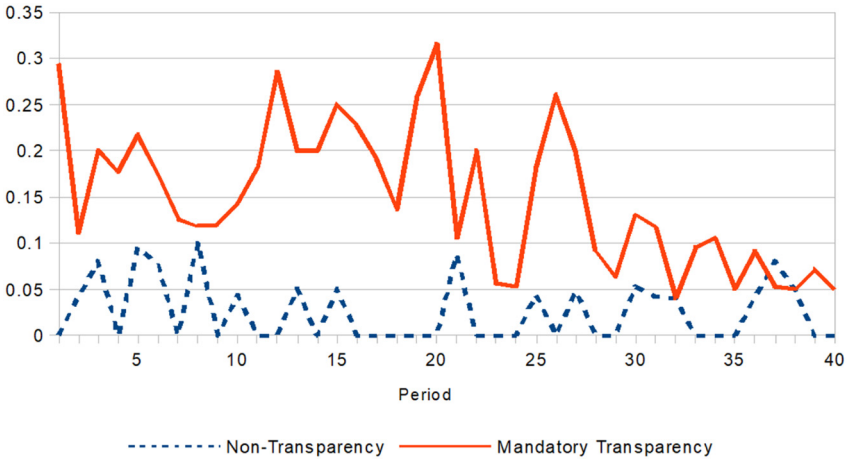

Fig. 2. Percentage of advisers who reject the payment and recommend r in state *R*.

equilibrium in the transparency treatment: (i) always r; (ii) choose r if $m \in \{r_A, r_R\}$ and l if $m = l$; (iii) choose r if $m \in \{r_A, l\}$ and l if $m = r_R$; and (iv) choose r if $m = r_A$ and l if $m \in \{l, r_R\}$.[27] In total, 72.7% of the decision makers use strategies that are consistent with the corrupt equilibrium.

There is only one strategy that is consistent with the honest equilibrium: choose r if $m = r_R$ and l if $m \in \{r_A, l\}$. As discussed before, the honest equilibrium requires a bias against an adviser who accepts the side payment. Indeed, we observe that a non-negligible fraction of decision makers are suspicious of advisers who accepted the side payment in the transparency treatments. 21% of decision makers use the honest equilibrium strategy; it is the second most popular strategy after the corrupt equilibrium strategy "choose r if $m \in \{r_A, r_R\}$ and l if $m = l$."

We next find which strategy is the optimal decision maker strategy in the transparency treatment given the empirical distribution of adviser behavior. To that aim, we compute the empirical posterior probability of state R given that (i) A accepts the payment and recommends r; (ii) A rejects the payment and recommends r; and (iii) A recommends l. We find that the optimal decision maker strategy is the corrupt equilibrium strategy, "choose r if $m \in \{r_A, r_R\}$ and l if $m = l$."

This finding might seemingly imply that it is suboptimal for decision makers to use the honest equilibrium strategy but there is an important caveat. The bias against advisers who accept the payment and the awareness of the advisers with respect to such a bias are likely the reasons why decision makers are better off with transparency. Arguably, the behavior of advisers is endogenous and shaped by the strategy of decision makers over the course of the experiment. Put differently, had the decision makers "always followed" the advisers, then we might have obtained a very different empirical distribution of adviser choices, and perhaps transparency would not have improved decision maker accuracy.

We now follow up on this line of reasoning and do a regression analysis in order to see whether (i) advisers who were previously (not) followed by the decision maker *due to* rejecting (accepting) the side payment are less likely to accept the payment; and (ii) advisers who were previously (not) followed by the decision maker *despite* accepting (rejecting) the side payment are more likely to accept the payment. We run a random effects panel probit regression of the probability A accepts the payment in the transparency treatment in round t (denoted by $Accept_t$) as a function of (i) the number of times

[27] The strategies except for the second one may seem unreasonable; however, IC and D1 do not rule out such strategies.

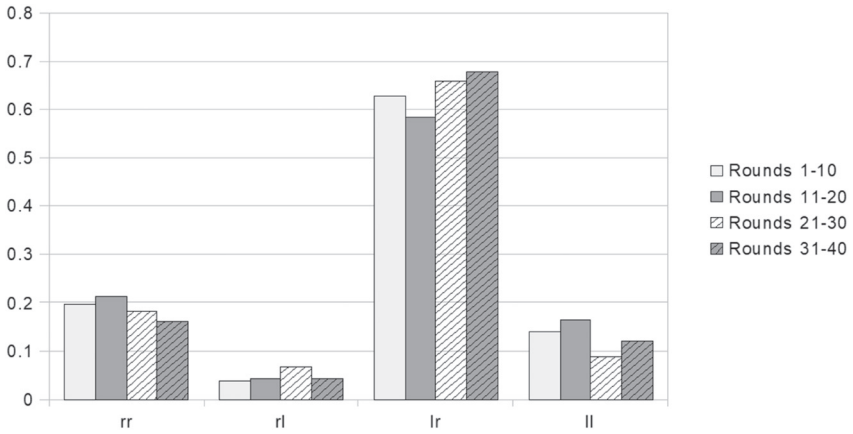

Fig. 3. Decision maker strategies in the non-transparency treatment.
Strategy xy indicates: x is chosen if A recommends l; y is chosen if A recommends r.

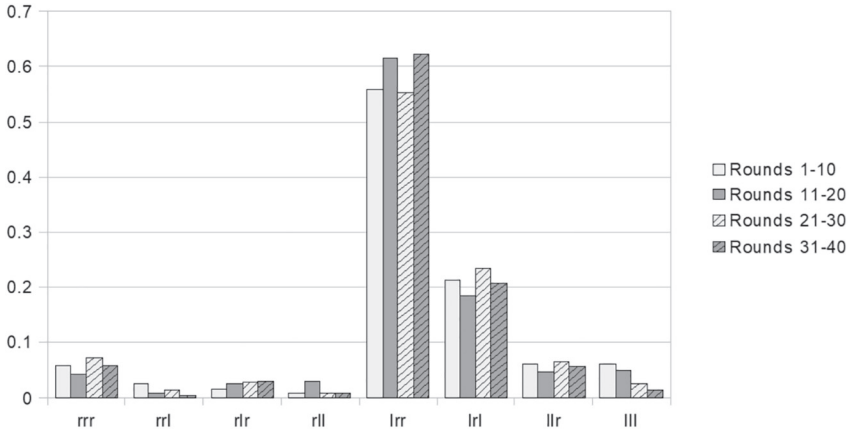

Fig. 4. Decision maker strategies in the transparency treatment.
Strategy xyz indicates: x is chosen if A recommends l; y is chosen if A rejects the payment and recommends r, z is chosen if A accepts the payment and recommends r.

A accepted the payment and was followed until t (denoted by FA); (ii) the number of times A accepted the payment and was *not* followed until t (denoted by NFA); (iii) the number of times A rejected the payment and was followed until t (denoted by FR); (iv) the lagged dependent variable; and (v) the round number (t). Thus, the panel model is given by[28]

$$Accept_{i,t} = \beta_0 + \beta_1 FA_{i,t-1} + \beta_2 NFA_{i,t-1} + \beta_3 FR_{i,t-1} + \beta_4 Accept_{i,t-1} + \beta_5 t + \varepsilon_{i,t} + \gamma_i.$$

We predict that $\beta_1 > 0$, $\beta_2 < 0$, and that $\beta_3 < 0$. The results of this regression are given in Table 4 under the column titled Model 1. The coefficient signs are as we predicted and the coefficients are significant. Model 2 is a variation on the same theme and involves the explanatory variables (i) the fraction of times the adviser was followed conditional on accepting the payment until t (denoted by $[P(F|A)]$); (ii) the fraction of times he was followed conditional on rejecting the payment until t (denoted by $P(F|R)$); (iii) the lagged dependent variable; and (iv) the round number (t). Thus, the panel model is given by

$$Accept_{i,t} = \beta_0 + \beta_1 [P(F|A)]_{i,t-1} + \beta_2 [P(F|R)]_{i,t-1} + \beta_3 Accept_{i,t-1} + \beta_4 t + \varepsilon_{i,t} + \gamma_i.$$

We predict that $\beta_1 > 0$ and $\beta_2 < 0$. The results of this regression are given in Table 4 under the column titled Model 2. The coefficient signs are as we predicted and the coefficients are significant.

These regression results support our notion that adviser behavior is endogenous and to a large extent shaped by the behavior of decision makers over the course of the experiment. These results also suggest that the percentage of adviser behavior consistent with the honest equilibrium declines over time due to the fact that the fraction of decision makers who are biased against and punish advisers who accept the side payment is not high enough. As discussed above, our role switching may have contributed to the time trend as well. Since decision makers in the first half of the transparency treatment were recommended the correct policy less than 1/3 of the time in state L, they may have been primed to accept the side payment frequently when they became advisers in the second half.

[28] To be more precise, the right hand side specifies the underlying latent propensity that $Accept_{i,t} = 1$.

Table 4
Random effects probit estimations of side payment acceptance.

	Coefficients	
Constant and independent variables	Model 1	Model 2
# Followed after Accept [*FA*]	0.109 (0.0114)***	
# Not followed after Accept [*NFA*]	−0.0708 (0.0372)*	
# Followed after Reject [*FR*]	−0.0971 (0.0307)***	
$Accept_{i,t-1}$	0.156 (0.121)	0.0481 (0.264)
% Followed after Accept [$P(F\|A)$]		1.195 (0.386)***
% Followed after Reject [$P(F\|R)$]		−0.735 (0.298)**
Round [t]	0.00837 (0.00879)	0.0140 (0.00912)
Constant	0.384 (0.151)**	0.260 (0.584)

Notes: (1) The dependent variable is the adviser's binary choice between accepting the payment (= 1) and rejecting (= 0). (2) The independent variables in Models 1 and 2 are explained in more detail in the main text. (3) Robust standard errors are given in parentheses. (4) * (**; ***) indicates significance at the 10% (5%; 1%) level. (5) Errors are clustered at the session level.

6.4. Voluntary transparency

Here, we give an overview of the key results of the voluntary-transparency treatment. A more thorough exposition can be found in Online Appendix A.

Voluntary transparency improved accuracy in state *L* (39.6%), but not as robustly as mandatory transparency. Accuracy in state *L* was 9.9% higher than in the non-transparency treatment, and 6.6% lower than in the mandatory-transparency treatment, the first difference being weakly significant while the latter is statistically insignificant. At 75.7%, accuracy in state *R* was almost identical to the other two treatments.

A majority of advisers chose to disclose their decision regarding the side payment: 68.6% chose to disclose in state *L* and 81.4% chose to do so in state *R*. Advisers who chose to be transparent and the decision makers with whom they were paired behaved very similarly to those under mandatory transparency. Advisers who choose not to be transparent behaved less honestly in state *L* than those in the non-transparency treatment. However, the decision makers with whom they were matched correctly recognized that a decision not to be transparent signaled dishonesty and less often followed subsequent advice, leading to accuracy rates similar to those in the non-transparency treatment. Thus, overall accuracy under voluntary transparency was intermediate between the other two treatments.

7. Conclusion

Our analysis has relevant implications for the debate on transparency. Our theoretical results show that in the environment we model, transparency is never harmful and can strictly improve the accuracy of decision making. Our experimental findings document that transparency makes decision makers better off in the state in which the adviser has an incentive to lie, and has no effect in the other state. With transparency, more advisers reject the side payment and recommend the correct policy in both states. While our paper documents positive effects of transparency, prior experimental research mostly produced bleak results. Previous studies modeled the conflicting interests of the adviser and the decision maker as being exogenous, but our results imply that the debate on transparency should take into account whether such conflicts should be modeled as being exogenous or endogenous. We believe that the latter is the appropriate approach as experts have in most cases the agency to accept or reject side payments, gifts and bribes.

There are three important qualifications to our findings. First, we find that the form of transparency may matter: while mandatory transparency clearly improves decision making relative to non-transparency (through the mechanism predicted by our theory) the evidence regarding the effect of voluntary transparency is weak. Second, the positive effects we document on adviser behavior weaken over time because many decision makers choose to follow the advice they receive even if the side payment is accepted, and advisers take advantage of this. Third, as mentioned earlier, in a pilot experiment run with equal side payments in each state, (mandatory) transparency was found to have no effect on decision-maker accuracy (we did not run sessions with voluntary transparency). Our intuition was that this was related to the honest equilibrium with the original parameters not being robust to the D1 criterion, which has been found to be relevant in previous work (see, for example, Banks et al., 1994). This guided the adjustments made for the experiment reported in this paper. It is tempting to see this as evidence that the D1 criterion is empirically relevant, but with only two sets of parameter values to compare, it is hard to generalize, and the different results may be due only to the increase in the side payment in one state, or the decrease in the other. It is not surprising that the degree to which transparency is effective depends on precise parameter values, as they affect factors such as the relative attractiveness of the honest equilibrium or the cost of disregarding truthful but paid-for advice. What we have shown in our experiment is that not only do there exist environments in which transparency is beneficial, but also that the mechanism predicted by our theoretical analysis can be empirically important, with a number of both advisers and decision makers playing strategies consistent with the honest equilibrium. We leave it to future research to determine the conditions under which transparency is more or less effective.

Online Appendices. Supplementary data for voluntary transparency condition, omitted theory and proofs, and experimental instructions

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jpubeco.2018.04.003.

References

Banks, J., Camerer, C., Porter, D., 1994. An experimental analysis of Nash refinements in signaling games. Games Econ. Behav. 6 (1), 1–31.

Behnk, S., Barreda-Tarrazona, I., García-Gallego, A., 2014. The role of ex post transparency in information transmission - An experiment. J. Econ. Behav. Organ. 101, 45–64.

Cain, D.M., Loewenstein, G., Moore, D.A., 2005. The dirt on coming clean: perverse effects of disclosing conflicts of interest. J. Leg. Stud. 34 (1), 1–25.

Cain, D.M., Loewenstein, G., Moore, D.A., 2011. When sunlight fails to disinfect: understanding the perverse effects of disclosing conflicts of interest. J. Consum. Res. 37 (5), 836–857.

Chen, D.L., Levonyan, V., Reinhart, S.E., Taksler, G.B., 2017. Mandatory disclosure: theory and evidence from industry-physician relationships. Working paper, Mimeo.

Cho, I.-K., Kreps, D.M., 1987. Signaling games and stable equilibria. Q. J. Econ. 102 (2), 179–221.

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

Delafontaine, R., 2015. Historians as Expert Judicial Witnesses in Tobacco Litigation: A Controversial Legal Practice. vol. 4. Springer.

Durbin, E., Iyer, G., 2009. Corruptible advice. Am. Econ. J. Microecon. 1 (2), 220–242.

Fanelli, D., 2009. How many scientists fabricate and falsify research? A systematic review and meta-analysis of survey data. Plos One 4 (5).e5738.

Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.

Gibson, R., Tanner, C., Wagner, A.F., 2013. Preferences for truthfulness: heterogeneity among and within individuals. Am. Econ. Rev. 103 (1), 532–548.

Gneezy, U., 2005. Deception: the role of consequences. Am. Econ. Rev. 95 (1), 384–394.

Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with orsee. J. Econ. Sci. Assoc. 1 (1), 114–125.

Hurkens, S., Kartik, N., 2009. Would i lie to you? On social preferences and lying aversion. Exp. Econ. 12 (2), 180–192.

Inderst, R., Ottaviani, M., 2012a. Competition through commissions and kickbacks. Am. Econ. Rev. 102 (2), 780–809.

Inderst, R., Ottaviani, M., 2012b. How (not) to pay for advice: a framework for consumer financial protection. J. Financ. Econ. 105 (2), 393–411.

Ismayilov, H., Potters, J., 2013. Disclosing advisor's interests neither hurts nor helps. J. Econ. Behav. Organ. 93, 314–320.
Kearns, C.E., Schmidt, L.A., Glantz, S.A., 2016. Sugar industry and coronary heart disease research: a historical analysis of internal industry documents. JAMA Int. Med. 176 (11), 1680–1685.
Kesselheim, A.S., Robertson, C.T., Myers, J.A., Rose, S.L., Gillet, V., Ross, K.M., Glynn, R.J., Joffe, S., Avorn, J., 2012. A randomized study of how physicians interpret research funding disclosures. N. Engl. J. Med. 367 (12), 1119–1127.
Koch, C., Schmidt, C., 2010. Disclosing conflicts of interest-do experience and reputation matter? Acc. Organ. Soc. 35 (1), 95–107.
Loewenstein, G., Cain, D.M., Sah, S., 2011. The limits of transparency: pitfalls and potential of disclosing conflicts of interest. Am. Econ. Rev. 101 (3), 423–428.
Loewenstein, G., Sah, S., Cain, D.M., 2012. The unintended consequences of conflict of interest disclosure. Jama 307 (7), 669–670.
Potters, J., Van Winden, F., 1992. Lobbying and asymmetric information. Public Choice 74 (3), 269–292.
Potters, J., Van Winden, F., 2000. Professionals and students in a lobbying experiment: professional rules of conduct and subject surrogacy. J. Econ. Behav. Organ. 43 (4), 499–522.
Rode, J., 2010. Truth and trust in communication: experiments on the effect of a competitive context. Games Econ. Behav. 68 (1), 325–338.
Sah, S., Loewenstein, G., 2014. Nothing to declare: mandatory and voluntary disclosure leads advisors to avoid conflicts of interest. Psychol. Sci. 25 (2), 575–584.
Sánchez-Pagés, S., Vorsatz, M., 2007. An experimental study of truth-telling in a sender-receiver game. Games Econ. Behav. 61 (1), 86–112.
Sobel, J., 1985. A theory of credibility. Rev. Econ. Stud. 52 (4), 557–573.
Stossel, T.P., 2005. Regulating academic-industrial research relationships-solving problems or stifling progress? N. Engl. J. Med. 353 (10), 1060.
Stossel, T.P., Stell, L.K., 2011. Time to 'walk the walk' about industry ties to enhance health. Nat. Med. 17 (4), 437–438.
Weber, M.A., 2009. Academic physicians confront a hostile world: the creation of acre. J. Clin. Hypertens. 11 (10), 533–536.