TI 2019-071/I
Tinbergen Institute Discussion Paper

Shooting the Messenger? Supply and Demand in Markets for Willful Ignorance

Shaul Shalvi[1]
Ivan Soraperra[1]
Joël van der Weele[1]
Marie Claire Villeval[2]

[1] University of Amsterdam, CREED
[2] University of Lyon, CNRS, GATE. IZA, Bonn

Shooting the Messenger?
Supply and Demand in Markets for Willful Ignorance*

Shaul Shalvi† Ivan Soraperra‡ Joël van der Weele§ Marie Claire Villeval¶

October 2, 2019

Abstract

We investigate the role of advisers in the transmission of ethically relevant information, a critical aspect of executive decision making in organizations. In our laboratory experiment, advisers are informed about the negative externalities associated with the decision-maker's choices and compete with other advisers. We find that advisers suppress about a quarter of "inconvenient" information. Suppression is not strategic, but based on the advisers' own preferences in the ethical dilemma. On the demand side, a substantial minority of decision makers avoid advisers who transmit inconvenient information (they "shoot the messenger"). Overall, by facilitating assortative matching, a competitive market for advisers efficiently caters to the demand for both information and information avoidance. Decision-makers are less likely to implement their preferred option when they are randomly matched to advisers and there is no scope for assortative matching.

Keywords: Self-deception, information avoidance, unethical behavior, experiment
JEL: C91, D82, D83, D91

*The authors are grateful to participants at the Beliefs workshop in Munich, the IMEBESS Conference in Utrecht, and the FELIS CORTEX workshop in Lyon for useful comments. The authors thank Bjorn Bartling for helpful discussions. They thank Q. Thevenet for programming, and J. Benistant, T. Garcia, R. Suchon, and C. Rimbaud for research assistance. This work has benefited from the financial support of the French National Research Agency (FELIS, ANR-14-CE280010-01), from IDEXLYON from Université de Lyon (project INDEPTH) within the Programme Investissements d'Avenir (ANR-16-IDEX-0005) and from the LABEX CORTEX (ANR-11-LABX-0042) of Université de Lyon, within the program Investissements d'Avenir (ANR-11-IDEX-007) operated by the French National Research Agency. It has also been supported by the Van Gogh PHC program. Joel van der Weele thanks financial support from the Dutch Science Association (NWO) via a personal VIDI grant (452-17-004).This research was received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (ERC-StG-637915).

†CREED, Univerity of Amsterdam. Email: S.Shalvi@uva.nl.
‡CREED, University of Amsterdam. Email: i.soraperra@uva.nl.
§CREED, University of Amsterdam. Email: J.J.vanderWeele@uva.nl.
¶Univ Lyon, CNRS, GATE. IZA, Bonn, Germany. villeval@gate.cnrs.fr.

1 Introduction

Executives often depend on advisers for their information. The relationship with these advisers thus becomes a key determinant of executive decisions: to the extent that advisers provide honest and transparent information about the consequences of different actions, they may encourage ethical conduct by the executives. However, they may also facilitate bad behavior by becoming "sycophants" or "yes-men" who withhold inconvenient information. Executives with weaker morals may further encourage such behavior by "shooting the messenger" of inconvenient truths. These dynamics are at the root of many corporate and political corruption scandals. For instance, investigating the behavior in the Watergate and Enron scandals Simon (2005) finds that "Deliberate ignorance and calculated ambiguity [...] were among the most salient and unattractive features of the wrongdoing in Watergate, and in more diluted form, they seem central to questions of laywer conduct in Enron."

Despite a large literature on the use and avoidance of inconvenient information in individual decisions,[1] we know little about the role of advisers in the supply of ethically relevant information. To investigate how supply and demand of ethically relevant advice affect ethical decision making we conduct a laboratory experiment. Subjects in the role of uninformed decision makers can choose to increase their profits, but at the risk of reducing the donation to a charity. Each decision maker is matched to an adviser, who earns money from each match and is informed about the consequences of each decision for the charity. The adviser can choose to relay potentially "inconvenient" information to the decision maker or can instead choose to send cheerful but irrelevant pictures.

To reflect different institutional arrangements, we study two schemes to match advisers and decision makers. The first is random matching, which reflects a system where informed advisers are independent and outside of the control of the executive. Since concerns for strategic reputation-building have little role in this setting, this independence potentially gives advisers the power to influence the executive in ways that best suit their own agenda. In the second scheme, the decision maker chooses the adviser and can switch adviser in each round of the game. This market arrangement shifts power to the executive who can now select either conscientious advisers or "yes-men", depending on her own inclinations. This gives advisers an incentive to tell decision makers what they want to hear.

We find that in the market setting, most advisers suppress ethically relevant information at

[1] A number of papers shows that people engage in "willful" or "strategic ignorance" of inconvenient information as an excuse of selfish behavior. The first studies demonstrating this behavior are Ehrich and Irwin (2005) and Dana et al. (2007), followed by fast growing number of replications and follow-ups (Nyborg, 2011; Conrads and Irlenbusch, 2013; Grossman, 2014; Feiler, 2014; Bartling et al., 2014; Exley, 2015; Kajackaite, 2015; van der Weele, 2013; Grossman and van der Weele, 2017). Related work shows how self-serving interpretations of ambiguity of risk and ambiguity increase selfishness in sharing decisions (Haisley and Weber, 2010; Di Tella et al., 2015; Exley, 2015; Garcia et al., 2019). Freddi (2019) provides evidence of information avoidance from the field.

least sometimes, and 25% of advisers suppress it most of the time. Suppression persists when we eliminate competition between advisers in the random matching treatment. Moreover, in both matching schemes, advice is highly correlated with adviser's own preferences in ethical dilemmas, suggesting it is not driven by strategic considerations. On the demand side, we find that when decision makers have the power to do so, about one-quarter "shoot the messenger" and avoid advisers that consistently provide informative messages. Overall, the degree of ignorance and selfish behavior in a market setting are indistinguishable from a situation where decision makers control their own information supply.

Our results show that on aggregate, the presence of advisers does not solve the problem of selective information acquisition that has been documented in the literature on individual decision making. In fact, markets for advice are remarkably efficient at allocating information and ignorance to those decision makers who demand it, allowing them to act mostly as they would when they are in full control of their information supply. The random matching arrangement reduces this efficiency by reducing information supply to information seekers and increasing it to avoiders. Interestingly, this does not lead to a decrease in ethical behavior, as decision makers who cannot avoid learning the consequences of their actions become more pro-social.

These insights are relevant in various applications. When it comes to organizational design, the results show that adviser independence affects the balance of power and limits the impact of the executive's preferences. However, it does not necessary lead to more ethical decision making, depending on the advisers' preferences. Furthermore, our results also relate to the sharing of information on social networks, which are designed to provide maximum freedom in deciding how to populate our news feed. Our study shows that such freedom likely leads to assortative matching, and does little to increase information quality about ethical consequences of one's actions. In fact, in a somewhat uncomfortable parallel to such platforms, irrelevant distractions make up a large part of the information shared in our experiment.

Our paper makes three main contributions to the existing literature. It is the first to consider markets for ethically relevant information and one of few studies to investigate the supply of ethical information. Coffman and Gotthard Real (2019) also consider advice in ethical dilemmas. Unlike in our study, advisers do not have an informational advantage and can only express their opinion. Being advised to be selfish deflects punishment by those who are hurt by the selfish actions, and decision makers act more selfishly in treatments where advisers are available. Lind et al. (2019) allow subjects in experimental ethical dilemmas to force information on decision makers even if they declined it, and show that this causes more decision makers to inform themselves. By contrast, by implementing competition between advisers we can analyze both the demand for biased information and the tendency to tell people what they want to hear.[2]

[2]There is a small literature on yes-men that studies the role of incentives in biased transmission of information within organizations (Prendergast, 1993). Opinion conformity with those of a manager has also been identified

Second, we contribute to a nascent literature on communication about social impact. For example, in Foerster and Van der Weele (2018b), a sender can share ethically relevant information about the impact of a donation decision with a receiver. Unlike in the present paper, *both* subjects can then make a donation. The authors investigate how image concerns affect the sender's decision to communicate. When the sender's own donation is made public, selfish senders excuse their actions by downplaying the impact of the donation, thus decreasing donations among the receivers of the information. Foerster and Van der Weele (2018a) and Bénabou et al. (2018) investigate the role of such exculpatory narratives theoretically.

Finally, we contribute to the emerging literature on group decisions and the dilution of responsibility. In particular, Falk and Szech (2013) show that more subjects consent to killing a mouse when there is joint responsibility. Bartling and Fischbacher (2012) show that people can partially avoid responsibility by delegating unkind actions to an intermediary. Studying the collaborative aspects of corrupt behavior, Weisel and Shalvi (2015) introduce complementarities in unethical behavior in a lying task, and show that lying is more prevalent in teams than in individual decision making. Kocher et al. (2018) find a strong dishonesty shift when individuals decide as group members that is driven by communication within groups. We contribute to this literature by studying whether the interactions between informed and uninformed players and assortative matching increase unethical decision making.

The remainder of this paper is organized as follows. Section 2 introduces our experimental design and procedures and presents our main behavioral conjectures. Section 3 presents our results. Section 4 discusses these results and concludes.

2 Design, Procedures and Conjectures

2.1 Experimental Design

The experimental design consists of three parts and two treatments. Only the third part differs across treatments. We first describe the first two parts before introducing our EXO and ENDO treatments. The instructions are available in Appendix A.[3]

2.1.1 Part 1: Elicitation of Social Preferences and the Demand for Ignorance

The first part is designed to elicit the social preferences of the participants under two successive information conditions. We inform participants that a €15 donation will be made by the experimenter to a charity, GiveDirectly, but depending on their decision, this donation can be

as a strategy of ingratiation for agents who have to compete for a promotion (Robin et al. (2014), see also Cummins and Nyman (2013)). Here we consider instead the moral domain and a setting where advisers and decision makers are independent.

[3]The experimental design was pre-registered at AsPredicted: https://aspredicted.org/blind.php?x=dr7743.

cancelled.[4] Participants have to make a first decision by choosing between two options under complete information. Option 1 pays them €9 and confirms the experimenter's donation to the charity, while option 2 pays them €15 but cancels the donation, introducing a moral dilemma. Before making their decision, participants can see a picture and a testimonial of a potential recipient of the donation taken from the website of GiveDirectly (see an example of picture in the instructions in Appendix A).

Then, participants have to make a second decision that is similar to the first, but under incomplete information, analogous to Dana et al. (2007). This decision gives us a measure of the demand for ignorance in a context in which there is no direct social interaction with others. The program determines randomly whether option 1 or option 2 cancels the donation, where either possibility is equally likely. Participants are not informed on the outcome of the random draw. However, before making their choice of option, they have to choose whether being informed about the consequences of their action for the charity. If they select "Beneficiary", they learn which option cancels the donation and their screen displays the picture and testimonial of a potential beneficiary before their choice of option. If they select "Cute animal", they remain uninformed: their screen displays an uninformative picture (a cute animal) and they will never learn the consequences of their action, neither before nor after their choice of option. The display of a cute animal is designed to capture a fun distraction of the kind we often encounter on the Internet, and to balance the use of recipient pictures when subjects receive information about the "Beneficiary".

After deciding on being informed or not and before choosing their option, participants are also asked to guess the number of other participants in the session selecting each type of picture. A correct guess pays €1. The objective is to let participants think about the extent to which other people prefer to remain ignorant.

2.1.2 Part 2: Role Familiarization

In the second part we familiarize participants with the two roles, advisers (called "senders" of information in the instructions) and decision-makers (called "receivers" in the instructions), that will characterize social interactions in part 3 through two incentivized choices. By introducing only some elements of the more complex environment that will be used in part 3, this part aims at helping participants to understand the two roles, regardless of the role they will eventually play in the following part. The decision is identical to the decision under uncertainty

[4] We informed participants that GiveDirectly (https://www.givedirectly.org) is a charity that transfers money to very poor families in developing countries and that this charity is rated as one of the 7 top charities in terms of cost-effectiveness by the charity evaluation site GiveWell, above many traditional charities in the world. We also distributed a document on the operating mode of GiveDirectly and displayed information from Wikipedia. We chose this charity because its website allows us to select pictures and testimonials of potential beneficiaries who have passed its screening.

made in part 1: option 1 pays €9 and option 2 pays €15 to the decision maker, and the program selects randomly for each decision maker which one of the two options cancels the donation to the charity, with a 0.5 probability for each option.

Participants are matched in groups with six other participants. First, all participants play in the role of an adviser. Advisers do not have to choose an option. With a 0.2 probability advisers are not informed which option cancels the donation and the picture of a cute animal is sent automatically to the decision maker. With a 0.8 probability they are informed. All the participants are asked to put themselves in the scenario in which advisers are informed and to decide whether to send or not information to a decision maker, both in the scenario that option 1 cancels the donation and in the scenario that option 2 cancels the donation. This gives us information on whether people prefer information or ignorance and whether they are willing to supply information or ignorance depending on whether news is "good" or "bad". (Throughout the paper, we will refer to news as "good" if option 1 cancels the donation since in that case, choosing option 2 maximizes the payoffs of both the decision maker and the charity. We will refer to news as "bad" if the more lucrative option 2 cancels the donation, since this generates an ethical trade-off between the decision maker and the charity.)

After making the choice as an adviser, all participants play in the role of a decision maker. Each participant is randomly matched with another player in the group of six. The decision maker's information depends on the choice of this other participant when he or she played in the role of an adviser. If this adviser has decided to share information, the decision maker screen indicates which option cancels the donation and displays the picture and testimonial of a potential recipient before the decision maker's choice of option. If the adviser has decided not to share information or in the case he or she was not informed (with a 0.2 probability), the screen displays the picture of a cute animal: the decision maker does not know which option cancels the donation and cannot see the picture of a potential beneficiary of the donation. When a decision maker can see the picture of a cute animal on the screen, he or she does not know whether he or she received this picture because the adviser selected it or because the adviser was uninformed himself or herself. It is common knowledge that if this part is selected for payment, players are paid based on their decision as a decision maker.

2.1.3 Part 3: Endogenous *vs.* Exogenous Matching Between Advisers and Decision Makers

In the third and main part of the experiment participants make choices in 25 identical periods, in one of two treatments EXO and ENDO. We first describe the features that are common to the two treatments. Participants remain matched with the same six other participants as in part 2. In each group, three participants are randomly assigned the role of advisers (identified

by a symbol: spade, diamond or club) while the four other participants are assigned the role of decision makers (identified by a letter and a number, R1 to R4). They keep the same role and identifier throughout the part. Decision makers have to choose one of the two same options as in the previous part; option 1 pays them €9 and option 2 €15. In each period, the program draws randomly and independently for each decision maker which option cancels the experimenter's donation to the charity; each option has a 0.5 probability to be selected. Decision makers are not informed of the outcome of this draw. By contrast, each adviser is informed of the consequences of the two options for the charity with a 0.8 probability for each of the four decision makers. In case he or she is informed, the adviser has to decide whether informing or not the decision maker (see screenshots in Appendix A). In the former case, the potential beneficiary's picture and testimonial are sent to the decision maker with information on which option cancels the donation. In the latter case or when the adviser is uninformed, the picture of a cute animal is sent to the decision maker with no information on which option cancels the donation, neither before nor after his or her choice of option.[5]

The difference between the EXO and the ENDO treatments lies in the matching process of advisers and decision makers. In both treatments, at the beginning of the period decision makers see a history box that displays a symbol for each type of information sent to him or her by each adviser in each of the *previous* periods. Symbols are either 'GD' for GiveDirectly—when the adviser sent information with the picture and testimonial of a potential recipient—or the symbol of an animal—if the adviser had no information or he/she received the information and decided to send the picture of the cute animal (see screenshots in A). The past choices of the advisers in the group are only visible to the decision makers, not to the other advisers.

In the EXO treatment before the decision maker chooses an option, he or she is randomly matched by the program with one of the advisers for the current period. He or she receives the information shared by this adviser for the current period and chooses one of the two options. In the ENDO treatment, after observing the history box, each decision maker has to select one of the advisers before choosing an option. Thus, in the ENDO treatment subjects can select which type of advisers they prefer, either those who are likely to share information truthfully or those who help them remain willingly ignorant. In both treatments advisers are paid €10 for each decision maker they are matched with, either exogenously in the EXO treatment or endogenously in the ENDO treatment. Advisers are not informed on the option chosen by the decision makers, regardless of whether they were matched with them or not.

[5]Note that when they are informed and before making their choices, advisors can see the pictures of the potential recipient and that of the cute animal, so that they cannot choose one or the other decision just to be able to observe such or such picture. Also, the same picture cannot be displayed in more than one period on a participant's screen.

2.2 Procedures

All sessions were conducted at GATE-Lab, Lyon, France. We ran 16 sessions (8 for the EXO treatment and 8 for the ENDO treatment). The 322 participants (161 in the EXO treatment and 161 in the ENDO treatment) are mainly students recruited from the local engineering, business and medical schools, using Hroot (Bock et al., 2014). 55.28% of the participants are females (57.14% in EXO and 53.42% in ENDO; two-sided Fisher's exact test, $p = 0.575$). The average age is 22.50 years (21.97 in EXO and 23.04 in ENDO; two-sided Mann-Whitney test, M-W hereafter, $p = 0.131$).[6] Table B.1 in Appendix B gives a summary of the sessions. The experiment was developed using z-Tree (Fischbacher, 2007).

Upon arrival, participants drew a tag from an opaque bag assigning them to a computer terminal in the lab. The instructions for each part were distributed and read aloud by the experimenter after completion of the previous part (see Appendix A). Together with the instructions of the first part participants received a description of GiveDirectly and of its operating mode taken from Wikipedia. Before playing the first and third parts, participants had to fill out a comprehension questionnaire. Questions were answered in private. At the end of part 3 a socio-demographic questionnaire was displayed on the participants' screen and then they received feedback on their earnings in the session.

The average duration of sessions was 75 minutes. At the end of the session the program randomly selected one of the 28 periods for payment (one of the two decisions in part 1, the decision as a decision maker in part 2 or one of the 25 periods in part 3). If a decision in part 1 or in part 2 was selected, participants received either €9 or €15, depending on their chosen option. If a period in part 3 was selected, the decision maker earned either €9 or €15, depending on the chosen option in that period; the adviser earned €10 for each decision maker he or she was matched with in that period (thus, the adviser minimally earned €0 if he or she was not matched to any decision maker in that period, and maximally earned €40 if he or she was matched with the four decision makers). GiveDirectly received a donation of €15 for each decision maker whose decision did not cancel the donation. The average payoff of the participants was €18.49 (standard deviation, S.D. hereafter, = 6.57), including a €5 show-up fee. Payments were made in cash, in a separate room and in private.

2.3 Behavioral Conjectures

Here we form a number of conjectures about the behavior we expect to observe. We distinguish between different types of decision makers depending on their motivation. *Selfish* decision makers are only motivated by the maximization of their individual payoff, they should choose option 2 in all treatments and should be indifferent between their information sources. By contrast,

[6] Except if specified otherwise, all the non-parametric tests reported in the paper are two-sided and take each individual as one unit of independent observation.

altruistic decision makers should choose option 2 if and only if they are informed that option 1 cancels the donation to avoid the risk of cancelling the donation by their decision. Indeed, having information about the consequences is a necessary condition to be altruistic, since in the absence of information, either option is equally likely to cancel the donation. Thus, in the ENDO treatment altruistic decision makers should select informative advisers, *i.e.*, those who in the past periods were the least likely to send animal pictures. Some decision makers may be described as *avoiders*; these agents want to choose the selfish option, but also want to maintain a positive self-image or avoid the guilt from being explicitly selfish. Remaining uninformed may serve as an excuse and help achieve both goals (Grossman and van der Weele, 2017). Thus, in the ENDO treatment such decision makers should try to match with an adviser who has sent less information in the previous periods. Furthermore, if in the current period the selected adviser has sent an informative picture, they may "shoot" this messenger by choosing another adviser in the future.

Advisers also have a number of interesting possible strategies. First, altruistic advisers who care about the charity would always want to convey news when they are bad in order to inform the decision maker of the possible trade-off. Non altruistic advisers will behave differently. If they anticipate a sufficient demand for ignorance in the ENDO treatment, they may *suppress* bad news, *i.e.*, send the uninformative picture of a cute animal, because they face competition for decision makers.[7] In the EXO treatment since senders are matched exogenously to decision makers, their beliefs about the decision makers' preferences should weigh much less than in the other treatment. There are no clear reasons to send anything else than the payoff-relevant information, but advisers may simply decide whether to send information or not based on their own preference for information or for ignorance (as revealed in part 1 of the experiment).

This analysis leads to the following behavioral conjectures.

Conjecture 1. *There is a demand for ignorance. A substantial fraction of the decision makers seek out advisers with uninformative messages.*

Conjecture 2. *There is a strategic supply of ignorance. More advisers suppress bad news in the ENDO than in the EXO treatment.*

Conjecture 3. *There is assortative matching in the ENDO treatment, with selfish or reluctant decision makers more likely to match with selfish or reluctant advisors.*

Conjecture 4. *In situations with an ethical trade-off, there is more selfish behavior by decision makers in the ENDO than in the EXO treatment.*

[7] Note however that suppressing *only* bad news may backfire, as a decision maker will learn that uninformative pictures signal bad news. Thus, an adviser who really wants to hide should also suppress some good news.

3 Results

We first give an overview of the type of news transmitted in both treatments and its impact. We then turn to analyze the demand and supply of news in the two treatments, and the matching of decision makers and advisors in the ENDO treatment. Finally, we look at the resulting distribution of ethical behavior in both treatments. "Good news" are messages that show no ethical trade-off, "bad news" are messages that show such a trade-off, and "no news" are cheerful animal pictures.

3.1 Overview of Advice Content and Use

The expected distribution of information available to advisers was 40% good news, 40% bad news and 20% no news.[8] Thus, if advisers transmitted all information or if decision makers selected only advisers who did so, this should be the distribution of information consumed by the decision makers in part 3. Figure 1 shows that the actual distribution of information observed by the decision makers differs starkly from this benchmark ($\chi^2(2) = 594.16$ and $\chi^2(2) = 668.54$ in ENDO and EXO, respectively).[9] With a prevalence of 40% in both treatments (precisely, 40% in ENDO and 42% in EXO), the consumption of no news is the most common in both treatments. By contrast, both good and bad news are observed by decision makers 30% of the time (precisely and respectively, 29% and 31% in ENDO; 29% and 29% in EXO). Interestingly, as we discuss in more detail below, the distribution of consumed information appears similar across treatments.

Does it matter what information decision makers consume? Figure 2 displays the fraction of choices for option 2 in part 3 by information condition and treatment, and shows that it matters a lot. In both treatments, decision makers choose almost systematically option 2 after no news (95% in ENDO and 92% in EXO) or good news (98% in ENDO and 96% in EXO). Since there is no ethical trade-off in these cases, this shows that subjects understand the choices in front of them. By contrast, when they get bad news, only 40% of decision makers choose selfishly, a fraction that is identical across treatments.

[8] The realized frequencies are: 38%, 41%, and 21% in ENDO; and 40%, 41%, and 19% in EXO. A goodness-of-fit test does not reject the null hypothesis that the realized distribution is the expected one in EXO ($\chi^2(2) = 2.21$; $p = 0.332$) but not in ENDO ($\chi^2(2) = 9.13$; $p = 0.010$).

[9] Repeating the test using the realized frequencies instead of the theoretical ones gives the same results.

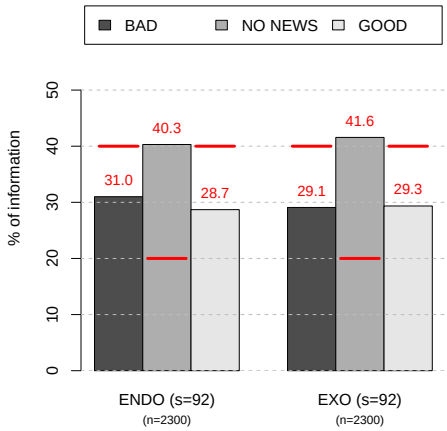

Figure 1: Information consumption

Notes: The figure displays the distribution of information observed by the decision makers in part 3, split by treatment. The horizontal lines show the distribution of information available to advisers. Labels below the bars indicate both the number of subjects (s) and the total number of choices (n).

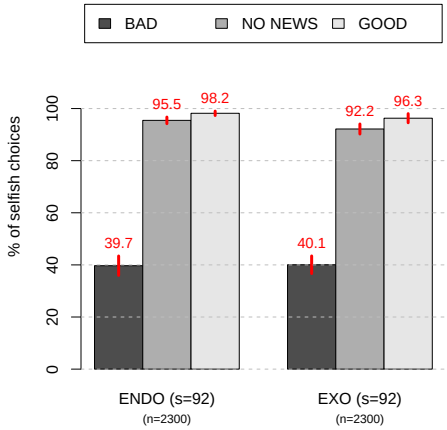

Figure 2: Information matters for the choice of option

Notes: The figure displays the fraction of times option 2 has been chosen by decision makers, split by treatment and information received. Vertical bars are standard errors based on a linear probability model with errors clustered at group level. Labels below the bars indicate both the number of subjects (s) and the total number of choices (n).

3.2 Demand for Advice

We now turn to the demand for advice in the ENDO treatment, where decision makers could choose an adviser. Our aim is to identify the strategies of decision makers when it comes to adviser selection, and connect it to their personal preferences as revealed in the individual part of the experiment (part 1). To facilitate the latter part of the exercise, we classify decision

makers' preferences in two-dimensional "types". The first dimension distinguishes "Selfish" and "Altruistic" decision makers, based on the choice of these decision makers in the first decision with full information about the ethical trade-off. The second dimension splits decision makers into (information) "Seekers" or "Avoiders" depending on whether they chose to receive information or not in the individual decision under uncertainty.

To understand demand for news, we consider two complementary approaches. The first is to examine the likelihood to choose different advisers depending on their profile of past advice, which is available to the decision maker. The second is to study the actual news consumed by a decision maker, and compare it to the general level of news available in the market. In Appendix B we provide a complementary analysis of the decision to switch adviser based on the content of the observed advice, which shows compatible results. In Appendix C we provide illustrations of the choice of adviser by individual decision makers.

Choice of adviser. Our first approach is to explain the choice of adviser based on each adviser's history of advice. To this end, we first rank the advisers according to the relative level of ignorance they provided to the decision maker in the previous 10 periods. We then ask how frequently the decision makers chose the adviser providing the highest, the intermediate, and a lowest level of ignorance. Note that this approach excludes the first 10 periods from the analysis, as advisers have not yet established a history.

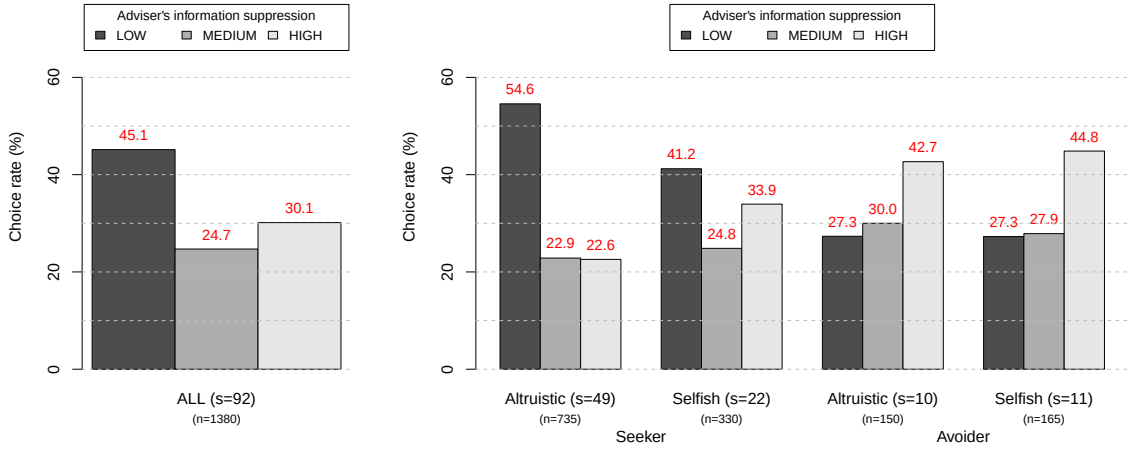

Figure 3: Switching advisers

Notes: The figure displays the frequency of choices of the three advisers in the ENDO treatment. Advisers are ranked (low, medium, and high) according to the relative level of ignorance they provided in the previous 10 periods. The left panel shows aggregate results. The right panel shows a split by decisions made by information seekers and information avoiders in part 1. Labels below the bars indicate both the number of subjects (s) and the total number of choices (n).

Figure 3 shows the frequencies with which different types of advisers are chosen. The left

panel shows that, on aggregate, the modal choice is the adviser that provides the highest level of information. The right panel shows a split by the "type" of decision makers, which reveals that the aggregate results hide a lot of heterogeneity. The largest group (Altruistic - Seekers) clearly penalizes non-informative advice: the most informative adviser is chosen 55% of the times and the least informative adviser is chosen 23% of the times. By contrast, a smaller group (Selfish - Avoiders) does the opposite: they choose the most informative adviser 27% of the times and the least informative adviser 45% of the times. Thus, decisions in the individual part of the experiment appear to have clear explanatory power for behavior in the interaction stage.

Table 1 evaluates these results statistically using multinomial logit models where the three alternatives are the advisers providing low, medium, and high ignorance. The individual specific explanatory variables include the dummies capturing the type of the decision maker, *i.e.*, Selfish-Altruistic and Avoider-Seeker, obtained from the choices in part 1. Both Model 1 and Model 2 include random effects at the individual level on the intercepts of the two equations. Compared to Model 1, Model 2 includes fixed effects at the group level using group dummies.

The estimates in Table 1 show the effect of the types on the odds ratio of Low versus Medium and of Low versus High ignorance, respectively. Moving from Altruistic to Selfish and from information Seeking to Avoiding significantly lowers the odds to choose the adviser that provides the lowest level of ignorance compared to the odds to choose the adviser that provides the intermediate level (Medium ignorance equation), and compared to the odds to choose the adviser that provides the highest level of ignorance (High ignorance equation). Results are robust to the inclusion of fixed effects at the group level. To help interpret these findings, Figure B.1 in Appendix B shows the predicted probability to choose each adviser for each decision maker. It shows that the effect of heterogeneity is mostly captured by shifting the probability mass from the Low ignorance to the High ignorance adviser, while the predicted probability to choose the medium ignorance adviser is about 20-25% and does not change much across types. Both the Selfish and Altruistic Avoiders show a significantly lower propensity to choose informative advisers.

Thus, our results show that subjects who avoided information in part 1 are likely to seek out uninformative advisers. For the Altruistic Avoiders, we can call this "strategic ignorance", as these subjects likely avoid bad news *in order* to avoid the cost of contributing to the charity. For the Selfish Avoiders, choosing uninformative advisers may not change their decision but may help them to avoid guilt from being confronted with the consequences of their actions.

Table 1: Probability to choose the adviser that provided the highest, intermediate and lowest level of ignorance.

	Model 1 Est. (S.E.)	Model 2 Est. (S.E.)
Medium ignorance		
(intercept)	-1.010 (0.109)***	-0.000 (0.379)
Selfish (DM)	0.349 (0.181)°	0.476 (0.200)*
Avoider (DM)	1.013 (0.247)***	1.089 (0.282)***
Selfish (DM) × Avoider (DM)	-0.675 (0.365)°	-0.844 (0.430)*
σ_M	0.779 (0.105)***	0.498 (0.127)***
High ignorance		
(intercept)	-1.037 (0.116)***	0.355 (0.405)
Selfish (DM)	0.848 (0.188)***	0.741 (8.218)***
Avoider (DM)	1.573 (0.253)***	1.647 (0.301)***
Selfish (DM) × Avoider (DM)	0.001 (0.385)	-0.917 (0.441)*
σ_H	1.560 (0.140)***	1.190 (0.139)***
Group dummies	NO	YES
Log-Likelihood	-1327.3	-1293.4
Number of observations	1380	1380
Number of subjects	92	92
Number of groups	23	23

Notes: These regressions are based on a multinomial logit model where the alternatives are the three advisers ordered by the amount of ignorance supplied in the previous 10 periods (the baseline alternative is the adviser that supplies the lowest level of ignorance). Individual specific variables are the dummies indicating the preferences of the decision makers. Both models include random effects at subject level on the intercepts. Model 2 includes group dummies. Regressions use data of the last 15 periods. *** ≤ 0.001; ** ≤ 0.01; * ≤ 0.05; ° ≤ 0.1.

News consumption. Our second method to classify decision makers' selection strategies is to compute how often a decision maker actually sees bad news, compared to the bad news that is available from advisers in the market. We focus on bad news, since this is the only news that matters from an ethical or efficiency perspective. We define the d-statistic, which is the fraction of bad states seen by the decision maker out of the average number of bad states reported to the decision maker (DM) by the three advisers he or she was matched with throughout the 25 periods of part 3:

$$d_i := \frac{\text{Number of bad states seen by DM } i}{\frac{1}{3}\sum_{j=1}^{3} \text{Bad states reported to the DM } i \text{ by adviser } j}.$$

Selecting an adviser at random will lead to $d \approx 1$. A decision maker who consistently selects an informative adviser will have $d > 1$, whereas selecting uninformative advisers will yield $d < 1$.

Figure 4 shows the distribution of the d-statistic. The left panel ranks all individual d-statistics by size, whereas the right panel shows the density distribution of the d-statistic. The left panel tells us that 40 out of 92 (43%) decision makers have a $d < 1$, and can be classified as information avoiders, while the rest consists of information seekers. Appendix C illustrates the demand for information by decision makers with different d-statistics. However, the d-statistic

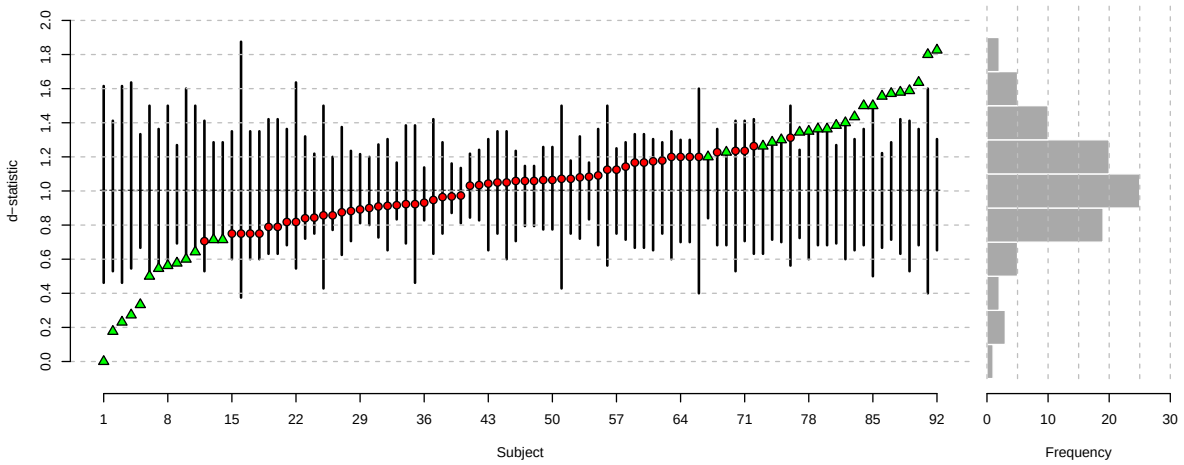

Figure 4: Information demand

Notes: The figure displays the distribution of the d-statistic across subjects in the ENDO treatment. A d-statistic equal to 1 corresponds to a random choice of advisers; a d-statistic higher (lower) than 1 corresponds to the selection of informative (uninformative) advisers. The left panel ranks the individual d-statistics by size and shows 90% confidence intervals based on simulations of random adviser choice. The right panel shows the overall distribution of the d-statistic.

is a noisy measure, since markets differ in the distribution of news, and hence in the possibility to become more or less informed. In an extreme case where all advisers transmit the same amount of news, the d-statistic will necessarily be 1, no matter what the news consumption is. We control for this by conducting simulations based on a decision maker who chooses randomly. This yields a distribution of d-statistics that we use to construct a 90% confidence interval.[10] By comparing the actual d-statistic to this confidence interval, we can classify with 90% confidence 13 decision makers as information avoiders (identifiable by triangles on the left hand side of the left panel), and 21 as information seekers (identifiable by triangles on the right hand side of the left panel), out of a total of 92. The percentage of avoiders in the set of classifiable decision makers is 38%, not too far from the 43% we found above.

Overall, we confirm the finding that a substantial minority of subjects appear to either avoid informative advisers or at the very least not seek them out. Coming back to our type classification of decision makers, the d-statistic correlates with the decision to avoid information in the first individual part of the experiment (Pearson $\rho = -0.202$, $p = 0.054$.) This leads to our first result that supports Conjecture 1.

[10]The procedure is as follows: (i) keeping fixed the advisers' behavior, we simulate the choice of each decision maker in each period under the assumption that he/she randomly selects one of the 3 advisers; (ii) given the simulated choices of the decision maker, we compute the implied d-statistic; (iii) we repeat the procedure 100.000 times. The confidence intervals are obtained by taking the 5th and 95th percentile of the simulated d-statistic over the 100.000 simulations. This interval captures the most likely values of the d-statistic under the assumption that the decision maker is neutral to the information received.

Result 1: *Using different metrics of information demand, we find that a majority of decision makers search for informative advisers. However, a substantial minority of decision makers do not:*

1. the *24% of subjects who avoid information in part 1 seek out non-informative advisers in part 3.*

2. *40% of subjects consume less bad news than the average amount available in the market.*

3.3 Supply of Advice

We now turn to the supply side of the market, where we investigate the extent of bad news suppression and its underlying motivations. As before, we focus on the suppression of bad news as this is the only news with ethical relevance.[11] We expect an increase in the suppression of bad news in the ENDO treatment, where advisers compete for clients and may try to satisfy the demand for information avoidance demonstrated in the previous section.

To measure bad news suppression, we define an adviser specific s-statistic, which is the fraction of bad news suppressed by an adviser divided by all bad states seen by that adviser in the 25 periods of part 3:

$$s_i := \frac{\text{Number of bad states suppressed by adviser } i}{\text{Number of bad states observed by adviser } i}$$

An adviser who suppresses all bad information will have an s-statistic of 1, while an adviser that transmits all bad news has an s- statistic of 0. This measure naturally excludes cases in which the adviser received no information. In our sample, we count 1% of advisers with a s-statistic of 1 and 29% with s-statistics of 0. Figure 5 shows the cumulative distribution of s-statistics over advisers in each treatment, which reveals several results. First, the distributions do not differ much by treatment. Indeed, a Kolmogorov-Smirnov test cannot reject equality of the distributions ($p = 0.248$). Second, about one third of advisers in each treatment transmits all bad news, and the large majority suppresses at least some news. Third, about 25% of advisers suppress more than half of the bad news they receive ($s > 0.5$). Appendix C gives individual examples to illustrate the supply of information by advisers with different s-statistics.

The absence of a treatment effect indicates that suppression is not (primarily) driven by strategic motives. To understand whether advisers' preferences can explain suppression behavior, Figure 6 shows correlations between the suppression rate of bad news in part 3 and the adviser's personal preferences toward both the donation and information as revealed in part 1

[11] We find substantial suppression of good news. However, this does not affect decision making, as illustrated in Figure 2. Suppressing good news can be understood as a complement to suppressing bad news, as advisers who only suppress bad news inadvertently reveal the news content through their suppression.

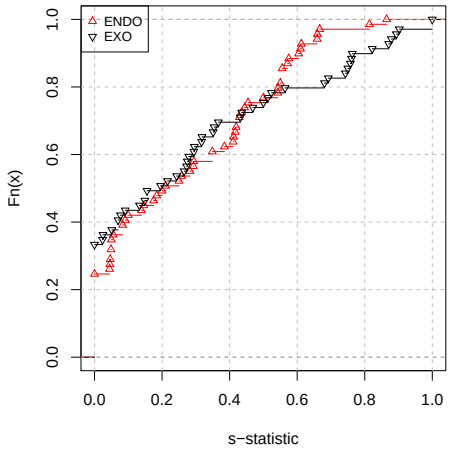

Figure 5: Suppression of bad news by advisers

Note: The figure displays the cumulative distribution of individual s-statistics by treatment.

of the experiment. The figure makes clear that both selfish advisers and advisers who avoided ethically relevant information for themselves are substantially less likely to transmit bad news, ranging from a suppression of 20% of bad news by Altruistic Seekers to 33% by Selfish Seekers, 42% by Altruistic Avoiders, up to 54% by Selfish Avoiders.

To further disentangle the importance of strategic and personal motivations by the advisers, Table 2 reports regressions with a dummy that takes value 1 when the adviser suppresses bad news in part 3 as the dependent variable. The explanatory variables in Model 1 have the EXO treatment as a baseline and include both dimensions of adviser types (Selfish-Altruistic, Avoider-Seeker) as well as interactions with the treatment. They also include the incentivized measure of beliefs about the fraction of Avoiders in their session that we elicited at the end of part 1. This is a good proxy for strategic motives: an adviser who believes that many people avoid information should be more likely to suppress information in the ENDO treatment, where there is competition for clients. The regressions include a time trend and controls for a series of individual characteristics: age and grade obtained at the final high school exam (Baccalaureat) measured as a deviation to the average of all subjects in the experiment, a dummy for gender and the reported number of participations in past experiments.

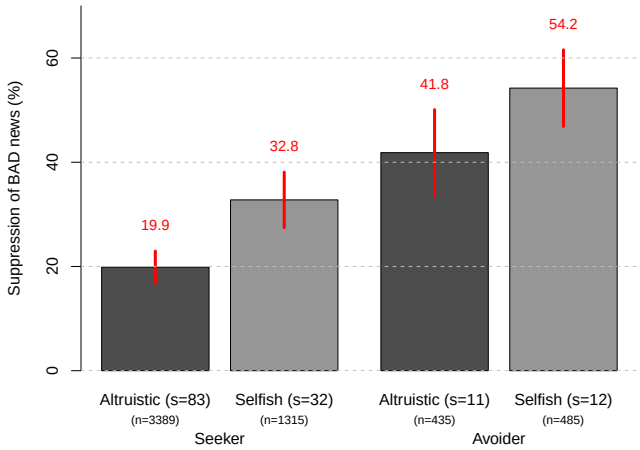

Figure 6: Adviser preferences and suppression of bad news

Notes: The figure shows the impact of adviser's preferences as revealed in part 1 of the experiment on the rate of suppression of bad news in part 3. Vertical bars show the standard errors. Standard errors are based on a linear probability model that clusters errors at group level and it is analogous to the ones reported in Table 2. Compared to the table, the model only includes the dummy Avoider, the dummy Selfish, and their interaction as explanatory variables.

Table 2: Suppression of bad news by advisers

	Model 1 Est. (S.E.)	Model 2 Est. (S.E.)
(Intercept)	0.075 (0.074)	0.071 (0.072)
Selfish	0.199 (0.097)*	0.222 (0.105)*
Avoider	0.134 (0.100)	0.186 (0.164)
Selfish × Avoider	— —	-0.134 (0.228)
Belief # ignorant	0.023 (0.010)*	0.023 (0.010)*
d(ENDO)	0.075 (0.066)	0.081 (0.067)
d(ENDO)×Selfish	-0.104 (0.112)	-0.164 (0.121)
d(ENDO)×Avoider	-0.045 (0.137)	-0.126 (0.192)
d(ENDO)× Selfish × Avoider	— —	0.314 (0.268)
d(ENDO)×Belief # ignorant	-0.019 (0.011)	-0.015 (0.011)
$Age - \overline{Age}$	-0.001 (0.003)	-0.001 (0.003)
d(Male)	-0.043 (0.042)	-0.040 (0.042)
$BAC - \overline{BAC}$	-0.012 (0.010)	-0.012 (0.010)
Number of past participations in exp.	-0.005 (0.016)	-0.006 (0.016)
Period dummies	YES	YES
Number of observations	5389	5389
Number of clusters	46	46

Notes: These regressions are based on linear probability models. The binary dependent variable is the adviser's choice to suppress bad news in part 3. Robust standard errors clustered at group level are in parentheses. *d* for dummy variables. "Belief # ignorant" is the subject's belief about the number of participants in their session that were willing to remain uninformed in part 1. Period dummies are included with period 1 as the reference category. * ≤ 0.05.

Model 1 confirms that Selfish and Avoider adviser types are more likely than others to suppress information in the EXO treatment, but the effect is significant only for the former type, at the 5% level. The effect of both types of preferences is less pronounced in the ENDO treatment (but not significantly so), in line with the idea that strategic motivations play a larger role in this treatment. The effect of the beliefs on the demand for ignorance has the expected sign and is significant. However, contrary to our expectations, beliefs are *less* correlated with suppression decisions in the ENDO treatment, but the difference is not statistically significant. Adding interaction terms between the adviser types (Selfish and Avoider) to the previous independent variables in Model 2 does not change the results, and the interaction coefficients themselves are very imprecisely estimated. This analysis leads to our second result that supports Conjecture 2.[12]

Result 2: *The majority of advisers do not transmit bad news to the decision makers and about 25% of advisers suppress more than half of the bad news they receive. There is no evidence of strategic behavior by advisers. Advisers' preferences have a strong impact on suppression decisions, especially in the EXO treatment.*

3.4 Assortative Matching

To complement our analysis of the demand and supply of ignorance, we next examine how different types of decision makers and advisers match with each other. Given our findings that both information supply and information demand are mainly driven by subjects' personal preferences, one would expect assortative matching to occur in the ENDO treatment, while it is ruled out by construction in the EXO treatment. To evaluate assortative matching, we use the two dimensional type classification explained above, based on part 1 choices. Table 3 shows the results of linear probability models where the dependent variable is equal to one if a decision maker matches with a Selfish adviser (Models 1-2) or an Avoider adviser (Models 3-4), and 0 otherwise. In addition to the decision maker's type, the independent variables include period dummies and individual characteristics of the decision maker.

Model 1 shows that being a Selfish decision maker hardly increases the probability of matching with a Selfish adviser. However, being an Avoider increases this probability by 17 percentage points, which is significant at the 5% level. On reflection, this pattern is not so surprising, given that being an Avoider is the strongest predictor of demand for ignorance in the market (see Table 1 and Figure 3), while being Selfish is the strongest predictor of supplying ignorance (see Table 2). If we add an interaction in column 2, we see that being Selfish Avoider increases the

[12] We also observe a correlation between part 1 choices and the decision to suppress bad news when considering part 2 data, where all the subjects had to choose whether to transmit information in a one-shot setting without competitive pressure. We observe that 18.1% of the information seekers suppress bad news compared to 42.9% of the information avoiders ($\chi^2(1) = 17.48$; $p < 0.001$). Similarly, 19.9% of the Altruists suppress bad news compared to 29.2% of the Selfish subjects ($\chi^2(1) = 3.50$; $p = 0.061$).

Table 3: Matching of types in ENDO

	Choosing a Selfish Adviser		Choosing an Avoider adviser	
	Model 1	Model 2	Model 3	Model 4
	Est. (S.E.)	Est. (S.E.)	Est. (S.E.)	Est. (S.E.)
(Intercept)	0.437 (0.086)***	0.442 (0.090)***	0.529 (0.103)***	0.550 (0.101)***
Selfish DM	0.018 (0.069)	0.007 (0.072)	-0.059 (0.054)	-0.096 (0.061)
Avoider DM	0.174 (0.074)*	0.153 (0.114)	0.210 (0.105)*	0.148 (0.111)
Selfish DM × Avoider DM	— —	0.044 (0.178)	— —	0.148 (0.132)
$Age - \overline{Age}$	-0.006 (0.003)*	-0.007 (0.003)*	-0.008 (0.002)***	-0.009 (0.002)***
d(Male)	-0.015 (0.088)	-0.017 (0.091)	-0.116 (0.049)*	-0.119 (0.047)*
$BAC - \overline{BAC}$	-0.006 (0.016)	-0.006 (0.016)	-0.017 (0.011)	-0.018 (0.011)
Number of past participations in exp.	0.032 (0.019)°	0.032 (0.019)°	0.031 (0.015)*	0.025 (0.015)°
Period dummies	YES	YES	YES	YES
Number of observations	1775	1775	1150	1150
Number of clusters	19	19	12	12

Notes: These regressions are based on linear probability models. The binary dependent variable is the decision maker's choice of a Selfish advisor in Models 1 and 2 or an advisor who is an Avoider in Models 3 and 4, in part 3. Robust standard errors clustered at group level are in parentheses. DM for decision maker; *d* for dummy variables. Period dummies are included with period 1 as the reference category. Regressions include only data from the groups where there was at least one adviser per type. *** ≤ 0.001; ** ≤ 0.01; * ≤ 0.05. Model 2 Wald test: $\beta = 0.153 + 0.044 = 0.197$, $p = 0.095$. Model 4 Wald test: $\beta = 0.148 + 0.148 = 0.296$, $p = 0.027$).

probability of a match with Selfish adviser by about 20 percentage points, which is marginally significant (see Wald test in Table notes). When we consider a match with an Avoider-type adviser as the outcome variable, we see very similar patterns. This analysis leads to our third result that supports Conjecture 3.

Result 3: *Endogenous matching leads Avoider-type decision makers to match significantly more with Selfish or Avoider-type advisers.*

3.5 Ethical Behavior With and Without Advisers

We now test our last conjecture by analyzing ethical decisions in part 3. To do so, we compare decisions in the two treatments in the "bad" state, where there was an ethical trade-off. Using the average amount of selfish decisions per individual as one observation (regardless of whether they were informed or uninformed), we find no evidence for a difference in the distribution of ethical decisions across the ENDO and the EXO treatments, with 62% of selfish decisions when option 2 cancels the donation in both treatments (Kolmogorov-Smirnov test, $p = 0.771$). Thus, at least on the surface, the level of adviser independence does not affect the amount of pro-social or selfish behavior.

We can also contrast decisions in the interaction part (part 3) with decisions under uncertainty in the individual part of the experiment, where decision makers had full control of their own information (the second decision in part 1). This comparison allows us to better

understand how the presence of advisers impacts ethical decisions.[13] The fraction of selfish decisions is 53% in the individual decision making part and 54% in both the EXO and ENDO treatments. This provides additional evidence that the presence of advisers does not lead to more or less ethical behavior.

To dig deeper, Figure 7 shows the results split by the decision makers' preferences for information, as revealed in part 1. This allows us to see how the loss of control over one's information supply affects subjects with different preferences for information. The first three bars in Figure 7 show individual decisions under full information (bar 1 for the first decision in part 1) and with uncertainty (bars 2 and 3 for Avoiders and information Seekers, respectively, in the second decision in part 1), where in the case of uncertainty we look only at cases where the state of the world was bad, *i.e.*, when there was an ethical dilemma. The fraction of selfish decisions increases substantially when subjects have the option to avoid information, in line with the results of Dana et al. (2007). Interestingly, this is true even for subjects who inform themselves, which may indicate some form of "moral licensing". The next two bars show the result for Avoiders and Seekers, respectively, in the ENDO treatment. Compared to the individual decision under uncertainty, we observe a small drop in selfishness among Avoiders in part 3, but overall there is still a large difference between Avoiders and Seekers of information. The last two bars reflect decisions in the EXO treatment. Selfish decisions now decline among Avoiders and increase among Seekers, yielding a much smaller gap.

These results show that when subjects can choose their advisers, their behavioral patterns are similar to the case where they are in control of their own information. In particular, the heterogeneity of adviser behavior shown in Figure 5 means that decision makers can obtain the information or non-information they prefer. By contrast, the loss of control in the EXO treatment means that Seekers may no longer get the information they need to distinguish unethical actions from ethical ones, and Avoiders will now regularly be confronted with information and may find harder to make selfish decisions with full knowledge of the consequences of selfish decisions. As a result, the behavioral patterns of the two groups come closer together and are in fact statistically indistinguishable, as we show below.

Table 4 reports a regression analysis of the determinants of unethical behavior under both matching protocols. We run linear probability models in which the dependent variable is the cancellation of the donation. Again, we focus on cases where the state was bad, and all advisers received information. The independent variables include the treatment, the preference types of the decision maker (Selfish and Avoider), as well as interaction terms between the preference type and the treatment, a time trend and the standard individual characteristics. As in the pre-

[13] To do so, we restrict the decisions in the interaction part to cases where all three advisers had information about the state of the world. This most closely mirrors the individual decision under uncertainty in part 1 where the information could be accessed by decision makers.

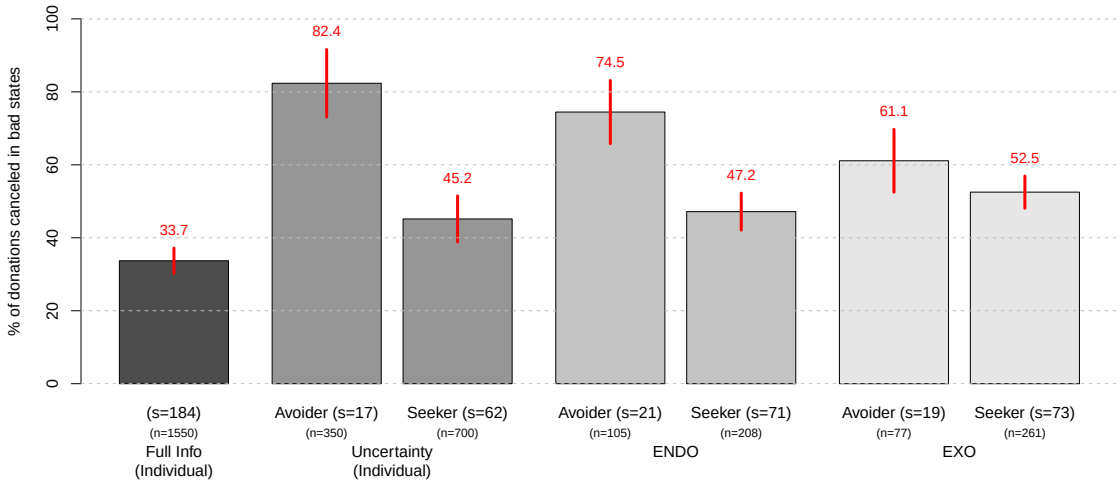

Figure 7: Selfish decisions in the bad state by condition

Notes: The figure displays the fraction of selfish decisions in the different experimental conditions. The first bar corresponds to the first decision in part 1; the second and third bars correspond to the second decision in part 1; the remaining bars correspond to decisions in part 3 with social interactions. Decision makers are split in Avoiders and Seekers of information, based on their decision in part 1 under uncertainty. Vertical lines are standard errors based on a linear probability model with standard errors clustered by groups for the four rightmost bars. To increase comparability with part 1 choices, the fraction of selfish decisions in ENDO and EXO are computed using cases where the state is bad and all advisers received information.

vious regression tables, the only difference between Model 2 and Model 1 is the inclusion of an interaction term between Selfish and Avoider types, itself interacted with the ENDO treatment.

Models 1 and 2 show that in both ENDO and EXO treatments Selfish decision makers are more likely than Altruists to make selfish decisions when they interact with advisers, and the correlation is highly significant. Consistent with Figure 7, information Avoiders do not behave significantly differently from information Seekers in EXO in any model. By contrast, in Model 1 a Wald test shows that being Avoider in the ENDO treatment ($\beta = 0.079 + 0.102$) increases the likelihood of making unethical decisions by 18 percentage points, which is marginally significant ($p = 0.094$). In Model 2, Wald tests show that this effect is mainly driven by the Avoiders who were Altruist in part 1. Indeed, in ENDO the effect of being Avoider has a positive and significant effect on the likelihood of cancelling the donation for an Altruistic decision maker ($\beta = 0.130 + 0.161 = 0.291$, $p = 0.046$), while the effect of being Avoider has no significant effect for a Selfish decision maker ($\beta = 0.130 - 0.166 + 0.161 - 0.073 = 0.291$, $p = 0.604$). This leads to our last result that does not support our Conjecture 4.

Result 4: *There is no evidence that the matching protocol affects aggregate unethical behavior. Aggregate selfishness in the presence of advisers does not differ from choices under uncertainty when decision makers control their information. By contrast, the matching protocol affects the relative impact of the preferences of decision makers for information: while this impact is*

Table 4: Determinants of unethical decisions

	Model 1 Est. (S.E.)	Model 2 Est. (S.E.)
(Intercept)	0.322 (0.090)***	0.311 (0.093)***
Selfish DM	0.493 (0.066)***	0.526 (0.075)***
Avoider DM	0.079 (0.075)	0.130 (0.100)
Selfish × Avoider DM	— —	-0.166 (0.148)
d(ENDO)	-0.051 (0.079)	-0.061 (0.081)
d(ENDO) × Selfish DM	0.009 (0.090)	0.043 (0.108)
d(ENDO) × Avoider DM	0.102 (0.129)	0.161 (0.173)
d(ENDO) × Selfish × Avoider DM	— —	-0.073 (0.195)
$Age - \overline{Age}$	-0.003 (0.002)	-0.003 (0.003)
d(Male)	0.056 (0.055)	0.055 (0.055)
$BAC - \overline{BAC}$	0.012 (0.013)	0.013 (0.013)
Number of past participations in exp.	-0.004 (0.016)	-0.005 (0.016)
Period dummies	YES	YES
Number of observations	1158	1158
Number of clusters	46	46

Notes: These regressions are based on linear probability models. The binary dependent variable is the cancellation of the donation in part 3. Robust standard errors clustered at group level are in parentheses. *d* for dummy variables. Data include only cases where the state is bad and all advisers received information. Period dummies are included with period 1 as the reference category. *** ≤ 0.001. Model 1 Wald test: effect of being an Avoider in the ENDO treatment $\beta = 0.079 + 0.102 = 0.181$, $p = 0.094$. Model 2 Wald tests: effect of being Avoider for Altruists in ENDO: $\beta = 0.130 + 0.161 = 0.291$, $p = 0.046$; effect of being Avoider for Selfish decision makers in ENDO $\beta = 0.130 - 0.166 + 0.161 - 0.073 = 0.291$, $p = 0.604$.

insignificant in the EXO treatment, altruistic decision makers who prefer avoiding inconvenient information are significantly more likely to cancel the donation in the ENDO treatment.

4 Conclusion

We study markets for the transmission of inconvenient information. On the demand side, we find substantial demand for both information and willful ignorance, in the form of irrelevant distractions. On the supply side, we observe substantial supply of ignorance by advisers, which appears to be driven by advisers' own preferences rather than strategic considerations when advisers have to compete with other advisers. Overall, markets for advice are remarkably efficient in catering to preferences for both information and ignorance. This is partly the result of assortative matching between advisers and decision makers, as decision makers search for the information supply that fits their own preferences. Thus, even though ethical behavior on aggregate does not differ between our matching protocols, decision makers' information preferences become more predictive of ethical behavior when advisers are chosen rather than exogenously assigned.

Our results have a number of applications. First, our setting can be seen as a stylized social media platform, somewhat akin to those outside the laboratory where many people obtain their information. It is therefore somewhat disconcerting that the most consumed information on

our lab platform consists of irrelevant distractions rather than ethically relevant information. Moreover, our results show similar dynamics of assortative matching as have been found on social media (e.g. Aiello et al., 2012). Our stylized setting demonstrates that as long as people can choose their own connections, they tend to select information from like-minded sources. As a result, they are unlikely to change their patterns of ethical behavior.

Second, our results matter for organizational design. The experiment shows how institutional details affect the balance of power in adviser-executive relations. Giving the executive the power to choose advisers will put a greater onus on the executive's character, while strengthening the independence of the bureaucracy does the opposite. However, having decision makers be guided by independent external advisers is not a guarantee for more ethical choices, as advisers suppress inconvenient information either to please the decision maker or to impose their own agenda. Given this tendency, organizations and individuals seeking to make well informed decisions may want to assure diversify of opinions among their advisors and solicit advice from a number of various sources.

References

Aiello, L. M., A. Barrat, R. Schifanella, C. Cattuto, B. Markines, and F. Menczer (2012): "Friendship prediction and homophily in social media," *ACM Transactions on the Web (TWEB)*, 6, 9.

Bartling, B., F. Engl, and R. A. Weber (2014): "Does willful ignorance deflect punishment? An experimental study," *European Economic Review*, 70, 512–524.

Bartling, B. and U. Fischbacher (2012): "Shifting the blame: On delegation and responsibility," *Review of Economic Studies*, 79, 67–87.

Bénabou, R., A. Falk, and J. Tirole (2018): "Narratives , Imperatives and Moral Reasoning," *NBER Working Paper 24798*, July.

Bock, O., I. Baetge, and A. Nicklisch (2014): "Hroot: Hamburg Registration and Organization Online Tool," *European Economic Review*, 71, 117–120.

Coffman, L. C. and A. Gotthard Real (2019): "Moral Perceptions of Advised Actions," *Management Science*, In press.

Conrads, J. and B. Irlenbusch (2013): "Strategic ignorance in ultimatum bargaining," *Journal of Economic Behavior & Organization*, 92, 104–115.

Cummins, J. G. and I. Nyman (2013): "Yes Men in Tournaments," *Journal of Institutional and Theoretical Economics*, 169, 621–659.

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

Di Tella, R., R. Perez-Truglia, A. Babino, and M. Sigman (2015): "Conveniently upset: Avoiding altruism by distorting beliefs about others' altruism," *American Economic Review*, 105, 3416–42.

Ehrich, K. R. and J. R. Irwin (2005): "Willful ignorance in the request for product attribute information," *Journal of Marketing Research*, XLii, 266–277.

Exley, C. L. (2015): "Excusing selfishness in charitable giving: The role of risk," *Review of Economic Studies*, 83, 587–628.

Falk, A. and N. Szech (2013): "Morals and markets," *Science (New York, N.Y.)*, 340, 707–11.

Feiler, L. (2014): "Testing Models of Information Avoidance with Binary Choice Dictator Games," *Journal of Economic Psychology*, 45, 253–267.

Fischbacher, U. (2007): "Z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10, 171–178.

Foerster, M. and J. J. Van der Weele (2018a): "Denial and Alarmism in Collective Action Problems," *Tinbergen Institute Discussion paper*, 019.

——— (2018b): "Persuasion , justification and the communication of social impact," *Tinbergen Institute Discussion paper*, 067.

Freddi, E. (2019): "Do People Avoid Morally Relevant Information? Evidence from the Refugee Crisis," *Review of Economics and Statistics*, In press.

Garcia, T., S. Massoni, and M. C. Villeval (2019): "Ambiguity and excuse-driven behavior in charitable giving," *SSRN*, 3283773.

Grossman, Z. (2014): "Strategic Ignorance and the Robustness of Social Preferences," *Management Science*, 60, 2659–2665.

Grossman, Z. and J. J. van der Weele (2017): "Self-Image and Willful Ignorance in Social Decisions," *Journal of the European Economic Association*, 15, 173–217.

Haisley, E. C. and R. A. Weber (2010): "Self-serving interpretations of ambiguity in other-regarding behavior," *Games and Economic Behavior*, 68, 614–625.

Kajackaite, A. (2015): "If I close my eyes, nobody will get hurt: The effect of ignorance on performance in a real-effort experiment," *Journal of Economic Behavior and Organization*, 116, 518–524.

Kocher, M. G., S. Schudy, and L. Spantig (2018): "I Lie? We Lie! Why? Experimental Evidence on a Dishonesty Shift in Groups," *Management Science*, 64, 3971–4470.

Lind, J. T., K. Nyborg, and A. Pauls (2019): "Save the planet or close your eyes? Testing strategic ignorance in a charity context," *Ecological Economics*, 161, 9–19.

Nyborg, K. (2011): "I dont want to hear about it: Rational ignorance among duty-oriented consumers," *Journal of Economic Behavior & Organization*, 79, 263–274.

Prendergast, C. (1993): "A theory of "yes men"," *American Economic Review*, 83, 757–770.

Robin, S., A. Rusinowska, and M. C. Villeval (2014): "Ingratiation: Experimental evidence," *European Economic Review*, 66, 16–38.

Simon, W. H. (2005): "Wrongs of Ignorance and Ambiguity: Lawyer Responsibility for Collective Misconduct," *Yale Journal on Regulation*, 22, 1–36.

van der Weele, J. J. (2013): "Inconvenient Truths : Determinants of Strategic Ignorance in Moral Dilemmas," *SSRN Electronic Journal*.

Weisel, O. and S. Shalvi (2015): "The collaborative roots of corruption," *Proceedings of the National Academy of Sciences*, 112, 10651–10656.

A Appendix: Instructions

These instructions were translated from French.

A.1 Instructions for the EXO treatment

Welcome to this experiment. Please switch off your mobile phone and refrain from communication with the other participants throughout the experiment, or we must exclude you from the experiment and from all payments. Please read the instructions carefully. Whenever you have a question, please raise your hand or press the red button on the side of your desk and we will come to your desk and answer to your question in private.

You will receive €5 for showing up on time. You can earn additional earnings based on your decisions and the decisions of other participants. The experiment consists of three parts that can include several periods. At the end of the session, the computer program will randomly select one of these periods, each with equal probability, and we will pay you according to your payoff in the selected period. Thus, you should think carefully when making each decision, as it could be the one that will be paid. At the end of the session, your total earnings will be paid to you in cash in a separate room.

For each participant, the experimenters have prepared a donation of €15 to a charity, GiveDirectly. GiveDirectly transfers money to very poor families in developing countries. This charity is rated as one of the 7 top charities in terms of cost-effectiveness by the charity evaluation site GiveWell, above many traditional charities in the world. Here is an excerpt from the website "GiveDirectly.org" presenting its objectives (we have also distributed a document on the operating mode of GiveDirectly and information from Wikipedia):

> *"We use mobile payments technology to send your donations to extremely poor families in the developing world in the most capital efficient way currently possible. $0.91 of your dollar ends up in the hands of the poor. Our model is setting the benchmark for philanthropic efficiency around the world. We strive to promote a new approach to philanthropy that uses constant experimentation and analytical rigor to understand the most impactful ways to achieve positive outcomes."*

During the session, we will show you pictures and testimonials of people who have passed the screening of GiveDirectly, and are potential recipients of the donations in this session. Their pictures and testimonials, translated into French, are taken verbatim from the website "GiveDirectly.org" and they may thus include typos.

The experimenters commit on honor to transfer the donations to GiveDirectly after the experiment. Note that the deontological rules of GATE-Lab do not allow deception of participants by the experimenters. So, all promised donations for the selected period at the end of the session will actually be sent to GiveDirectly. If you want more information about the transfer, please contact an experimenter after the session.

However, as we explain below, your choices may lead to a cancellation of the donation prepared by the experimenters, in which case GiveDirectly will not receive a donation for your participation.

The instructions for the first part follow below. The instructions for the next parts will be distributed after all participants have completed each part.

Part 1

In this part you will make two decisions. In each of these decisions, you are asked to choose between "Option 1" and "Option 2". Both options affect your own payoffs and the donation to GiveDirectly.

For **Decision one** you will see on your screen before your choice the picture and testimonial of a potential recipient of the donation, who has passed the screening by GiveDirectly, as illustrated in the screenshot below. Choosing "Option 1" will result in €9 for yourself and will *not cancel* the donation of €15 by the experimenters to GiveDirectly. Choosing "Option 2" will result in €15 for yourself, and will *cancel* the donation to GiveDirectly.

In **Decision two** as in decision one, you can choose between "Option 1", which will result in €9 for yourself and "Option 2", which will result in €15 for yourself. The difference with the first decision is that the program determines randomly which one of the two options will result in a cancellation of the donation of €15 to Give Directly. With 50 chances out of 100, choosing "Option 2" cancels the donation while "Option 1" confirms the donation, just like in decision 1. With 50 chances out of 100 the situation is reversed, so choosing "Option

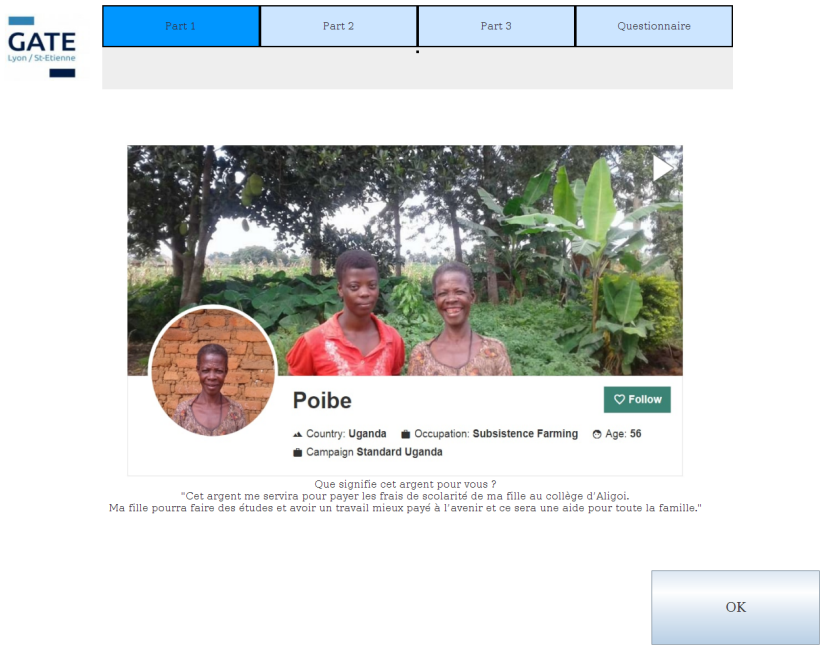

1" cancels the donation and "Option 2" confirms the donation. You are not informed which situation is chosen by the program, and the consequences for GiveDirectly are replaced by "???".

Before choosing between "Option 1" and "Option 2", you have to choose between two types of information.

- You can choose "Recipient". This means that before choosing between "Option 1" and "Option 2", you will learn which situation was chosen by the computer, and the "???" will be replaced with information about the consequences for GiveDirectly. Furthermore, like in decision 1, your screen will display the picture and testimonial of a potential recipient before your choice of option.
- Or you can choose "Cute animal". This means that your screen will display the picture of a cute animal, as illustrated in the screenshot below. You will not learn which situation was selected by the computer, neither before nor after your choice of option.

After making this choice, we will inform you about the number of participants in this session and ask you to guess the number of participants who have chosen "Recipient" and the number of participants who have chosen "Cute animal". Regardless of whether this period is selected or not for payment at the end of the session, you will earn 1 euro if your guess is correct, and 0 euro otherwise. Therefore, you should try to guess as accurately as possible. You will be informed on whether your guess is correct at the end of the session.

After you have chosen between "Recipient" and "Cute animal" and reported your guess, you will have to choose between "Option 1" and "Option 2". Your earlier choice between "Recipient" or 'Cute animal" determines the information you see on your screen before making your choice.

Summary of the Decisions

1. In Decision 1, you choose between Option 1 and Option 2.
2. In Decision 2, the program randomly selects which one of the two options cancels the donation.
3. You choose between the sets of information "Recipient" or "Cute animal".
4. You report your guess about the numbers of other participants in the session who chose 'Recipient" or "Cute animal".
5. Your screen displays the information you chose in step 3, and you choose between Option 1 and Option 2.

Please read again these instructions. If you have any questions, please raise your hand or press the red button. A comprehension questionnaire will be displayed on your screen.

OK

Part 2

In this part, you are randomly matched with 6 other participants to form a group of 7. There are two roles: Receivers and Senders. Receivers and Senders refer not to donations but to pictures and information, as explained below. All the participants in the group will first make decisions as Senders. Then, all of them will make a decision as Receivers. We first describe each role before explaining decision-making.

Choice of the Receiver

The Receiver has to choose between "Option 1" and "Option 2". The consequences from this decision are the same as in the second decision of part 1:

- "Option 1" results in €9 for the Receiver and "Option 2" results in €15 for the Receiver.
- The program picks randomly which one of the two options cancels the donation to GiveDirectly. Each option has 50 chances out of 100 to be picked.

The program randomly determines the consequences of each option independently for each Receiver. Thus, these consequences can differ across Receivers. Before making a choice between Option 1 and Option 2, the Receiver is not informed of the consequences of each option for GiveDirectly. However, s/he can obtain information from the Sender, as we now describe.

Choice of the Sender

With 80% chance, the Sender learns which one of the two options cancels the donation. With 20% chance the Sender does not learn the consequences of each option.

- If the Sender does not learn the consequences of each option for GiveDirectly, the program displays automatically the picture of a cute animal on the Receiver's screen before s/he makes his/her choice. The Receiver is not informed on the consequences of this option for GiveDirectly.
- If the Sender learns the consequences of each option for GiveDirectly, s/he has to choose between two types of information for the Receiver. If s/he chooses "Recipient", the Receiver will learn which one of the two options cancels the donation before choosing an option, and s/he will see the picture and the testimonial of a potential recipient of the donation. If the Sender chooses "Cute animal", the Receiver will see the picture of a cute animal, but not the consequences for GiveDirectly, neither before nor after the choice of option.

Decision Making in Part 2

In this part, all the participants first make two decisions in the role of a Sender in the case they are informed about the consequences of each option for GiveDirectly. Precisely, as a Sender, you have to choose between two types of information for the Receiver, either "Recipient" or 'Cute animal":

- in the case you learn that the donation to GiveDirectly is cancelled after "Option 1" , but not after "Option 2";
- and in the case you learn that the donation to GiveDirectly is cancelled after "Option 2", but not after "Option 1".

Then, all the participants will make a decision as Receivers. As a Receiver, you will have to choose between "Option 1" and "Option 2". Before you make your choice, the computer will randomly determine which option cancels the donation. It will also randomly pair you with a Sender in your group. The choice of the Sender between "Recipient" or 'Cute animal" determines the information you have about the consequences of each option.

Summary of the Decisions

1. You first decide as a Sender which picture and information to share if you are informed of the consequences of each option for Give Directly.
2. You are next a Receiver. You are randomly matched with a Sender.
3. You obtain the information chosen by the sender, "Recipient" or 'Cute animal".
4. You choose between Option 1 and Option 2.
5. You are paid based on your choice as a Receiver in case this part is selected for payment.

Please read again these instructions. If you have any questions, please raise your hand or press the red button.

Part 3

In this part, you are still matched with the same 6 other participants as in part 2. But now, participants are randomly assigned to one of the roles and will be identified with an ID. There are four Receivers and their IDs are R1, R2, R3, and R4. There are three Senders and their IDs are symbols (spade, diamond, club). We will communicate your role and your ID on your screen at the beginning of this part. This part has 25 identical periods and you will keep the same role and the same ID throughout this part. We now describe each of these periods.

Choice of the Receiver

In each period, the Receiver chooses between "Option 1" and "Option 2". The consequences from this decision are the same as before:

- "Option 1" results in €9 for the Receiver and "Option 2" results in €15 for the Receiver.
- The program picks randomly which one of the two options cancels the donation to GiveDirectly. Each option has 50% chance to be picked.

In each period, the program randomly determines the consequences of each option for GiveDirectly, independently for each Receiver. Thus, these consequences can differ across periods and across Receivers.

The Receiver is not informed about the consequences of each option for GiveDirectly. Before choosing between Option 1 and Option 2, s/he can receive information from the Sender.

Choice of the Sender

With 80% chance, the Sender learns the consequences of each option for GiveDirectly chosen by the program for each Receiver for the current period. If the Sender learns the consequences, s/he has to decide which set of information to share with the Receiver. As before, if s/he chooses "Recipient", i) the Receiver is informed which option cancels the donation before choosing an option, and ii) the picture of a potential recipient with his/her testimonial is displayed. If the Sender chooses "cute animal" the picture of a cute animal is displayed on the Receiver's screen and the Receiver does not learn which option cancels the donation, neither before nor after the

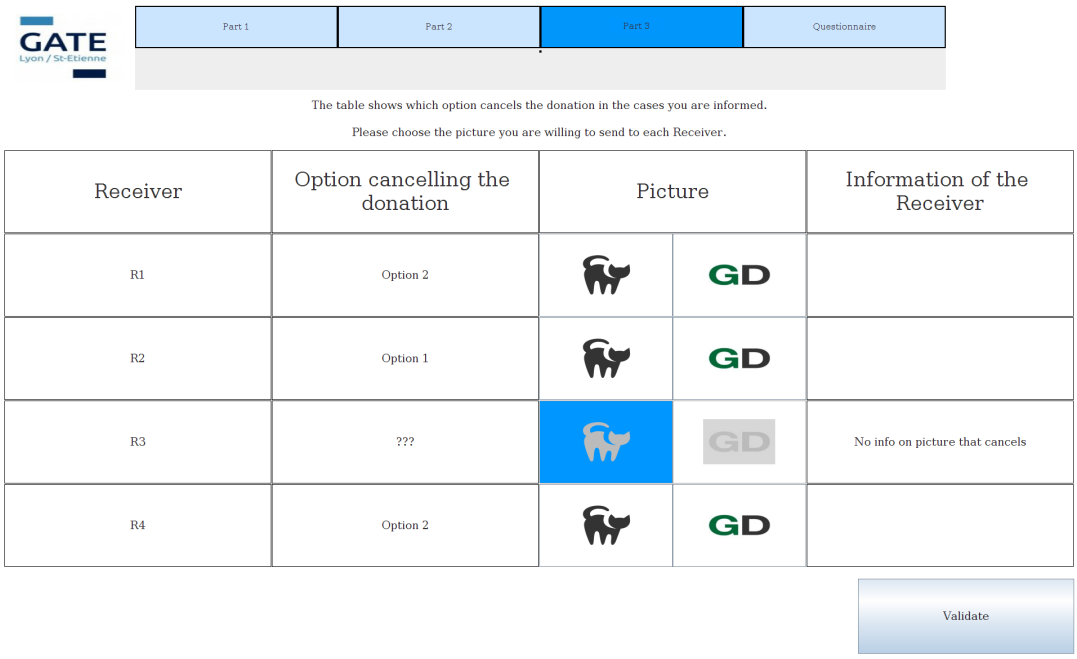

choice of option.

With 20% chance the Sender does not learn the consequences of each option and the picture of a cute animal is displayed automatically on the Receiver' s screen. The Receiver does not know whether the Sender has been informed or not.

The decision of the Sender is illustrated in the screenshot below. The first column of the table shows the ID of the Receivers (i.e. R1, R2, R3 or R4). The second column indicates for each Receiver, which option cancels the donation. The "???" sign indicates that the Sender did not receive information for participant R3; in this case, the sign of a cute animal is automatically pre-selected with no action from the Sender. In the next column, when informed, the Sender has to choose between "Recipient" and "Cute animal". In this example, the Sender makes three decisions, as s/he has information about the consequences of each option for three out of four Receivers. Once the Sender has made his/her decisions, the last column of the table indicates which information will be displayed on the screen of the Receiver.

As we explain now, the Sender can earn €10 for each Receiver to whom s/he has been randomly matched by the program in that period.

The Receiver is Matched with a Sender

Before the Receiver chooses an option, s/he is RANDOMLY matched with one of the Senders (spade, diamond, or club) for the current period. The screenshot below reflects the screen the Receiver will see in the experiment. The example shows period 5. The first column shows the ID of each Sender. For each of the past periods, the screen shows which set of information each Sender shared with the Receiver (indicated by the symbol of an animal or the symbol GD for a recipient of GiveDirectly). The symbol of an animal reflects *either* that the Sender had no information, *or* that the Sender received the information and decided to share this picture; the Receiver cannot distinguish between these possibilities. The past choices of the Senders in the group are only visible to the Receivers, not to the other Senders. In this example, symbols have been chosen randomly.

Before being randomly matched by the program to a Sender, the Receiver only knows the past choices of the Senders but not those for the current period. The Receiver's information ("Recipient" or "Cute animal") is determined by the randomly matched sender's choice for the current period. The Senders are not informed of the choices of the Receivers.

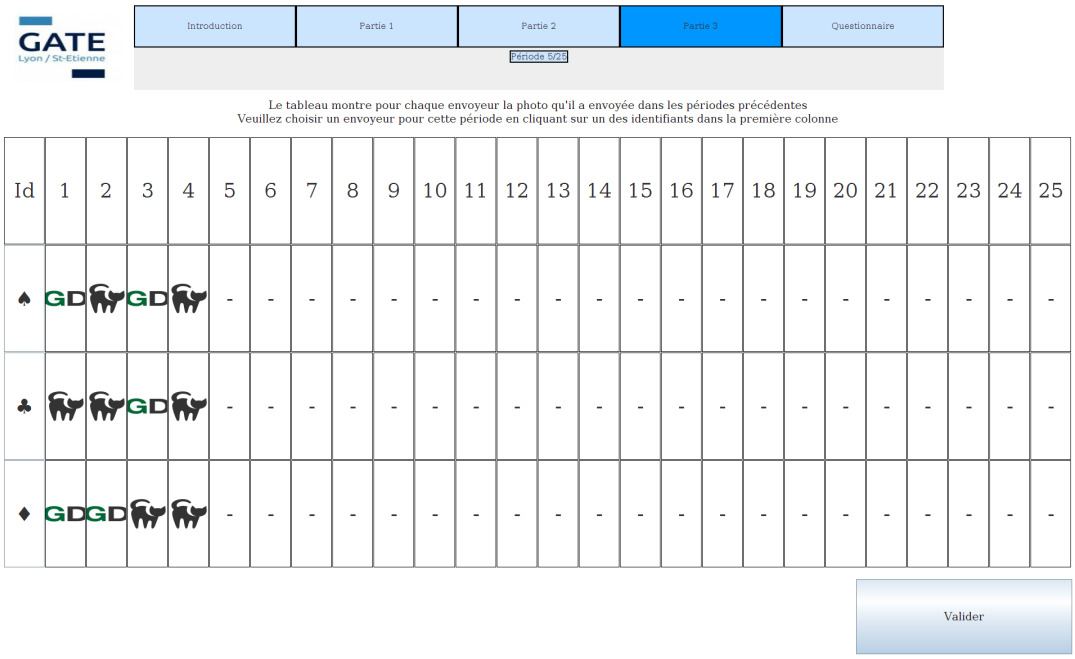

Summary of the Decisions

1. For each Receiver, the program randomly selects which option cancels the donation.
2. With 80% chance, the Sender learns the consequences of each option. If s/he learns the consequences, the Sender has to choose a set of information to share with the Receiver ("Recipient" or "Cute animal").
3. The Receiver sees information about the senders' choices in earlier periods. S/he is randomly matched by the program to one of the Senders. The Receiver's information is determined by the choice of that sender ("Recipien" or "Cute animal") in the current period.
4. The Receiver decides between Option 1 and Option 2.
5. The same procedure applies for the 25 periods.

Summary of Earnings

The program randomly selects one of the 28 periods for payment (the two decisions in part 1, the decision as a Receiver in part 2 and the 25 periods in part 3). If a decision in part 1 or in part 2 is selected for payment, you will receive either €9 or €15, depending on whether you chose option Option 1 or Option 2 in that period. If a decision in part 3 is selected, payoffs for each player are as follows:

- **The Receiver** earns either €9 or €15, depending on the chosen option in that period.
- **The Sender** earns €10 for each Receiver to whom s/he has been randomly matched by the program in that period. Thus, the Sender minimally earns €0 if s/he has not been matched to any Receiver, and maximally earns €40 if s/he has been matched to the four Receivers.
- Finally, **GiveDirectly** will receive a donation of €15 for any choice in the selected period that does not cancel the donation.

End of the Session

At the end of part 3 a questionnaire will be displayed on your screen and then you will receive a feedback on your earnings in the session. On invitation of an experimenter, you will move into the payment room with your pre-filled receipt of payment and your computer tag.

Please read again these instructions. If you have any questions, please raise your hand or press the red button. A comprehension questionnaire will be displayed on your screen.

A.2 Instructions for the ENDO treatment

Welcome to this experiment. Please switch off your mobile phone and refrain from communication with the other participants throughout the experiment, or we must exclude you from the experiment and from all payments. Please read the instructions carefully. Whenever you have a question, please raise your hand or press the red button on the side of your desk and we will come to your desk and answer to your question in private.

You will receive €5 for showing up on time. You can earn additional earnings based on your decisions and the decisions of other participants. The experiment consists of three parts that can include several periods. At the end of the session, the computer program will randomly select one of these periods, each with equal probability, and we will pay you according to your payoff in the selected period. Thus, you should think carefully when making each decision, as it could be the one that will be paid. At the end of the session, your total earnings will be paid to you in cash in a separate room.

For each participant, the experimenters have prepared a donation of €15 to a charity, GiveDirectly. GiveDirectly transfers money to very poor families in developing countries. This charity is rated as one of the 7 top charities in terms of cost-effectiveness by the charity evaluation site GiveWell, above many traditional charities in the world. Here is an excerpt from the website "GiveDirectly.org" presenting its objectives (we have also distributed a document on the operating mode of GiveDirectly and information from Wikipedia):

> *"We use mobile payments technology to send your donations to extremely poor families in the developing world in the most capital efficient way currently possible. $0.91 of your dollar ends up in the hands of the poor. Our model is setting the benchmark for philanthropic efficiency around the world. We strive to promote a new approach to philanthropy that uses constant experimentation and analytical rigor to understand the most impactful ways to achieve positive outcomes."*

During the session, we will show you pictures and testimonials of people who have passed the screening of GiveDirectly, and are potential recipients of the donations in this session. Their pictures and testimonials, translated into French, are taken verbatim from the website "GiveDirectly.org" and they may thus include typos.

The experimenters commit on honor to transfer the donations to GiveDirectly after the experiment. Note that the deontological rules of GATE-Lab do not allow deception of participants by the experimenters. So, all promised donations for the selected period at the end of the session will actually be sent to GiveDirectly. If you want more information about the transfer, please contact an experimenter after the session.

However, as we explain below, your choices may lead to a cancellation of the donation prepared by the experimenters, in which case GiveDirectly will not receive a donation for your participation.

The instructions for the first part follow below. The instructions for the next parts will be distributed after all participants have completed each part.

Part 1

In this part you will make two decisions. In each of these decisions, you are asked to choose between "Option 1" and "Option 2". Both options affect your own payoffs and the donation to GiveDirectly.

For **Decision one** you will see on your screen before your choice the picture and testimonial of a potential recipient of the donation, who has passed the screening by GiveDirectly, as illustrated in the screenshot below. Choosing "Option 1" will result in €9 for yourself and will *not cancel* the donation of €15 by the experimenters to GiveDirectly. Choosing "Option 2" will result in €15 for yourself, and will *cancel* the donation to GiveDirectly.

In **Decision two** as in decision one, you can choose between "Option 1", which will result in €9 for yourself and "Option 2", which will result in €15 for yourself. The difference with the first decision is that the program determines randomly which one of the two options will result in a cancellation of the donation of €15 to Give Directly. With 50 chances out of 100, choosing "Option 2" cancels the donation while "Option 1" confirms the donation, just like in decision 1. With 50 chances out of 100 the situation is reversed, so choosing "Option

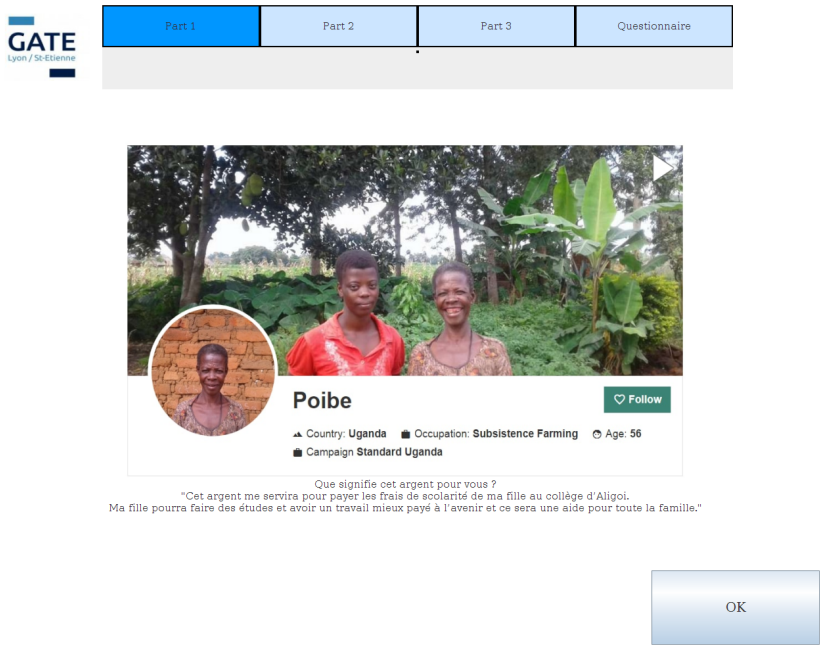

1" cancels the donation and "Option 2" confirms the donation. You are not informed which situation is chosen by the program, and the consequences for GiveDirectly are replaced by "???".

Before choosing between "Option 1" and "Option 2", you have to choose between two types of information.

- You can choose "Recipient". This means that before choosing between "Option 1" and "Option 2", you will learn which situation was chosen by the computer, and the "???" will be replaced with information about the consequences for GiveDirectly. Furthermore, like in decision 1, your screen will display the picture and testimonial of a potential recipient before your choice of option.
- Or you can choose "Cute animal". This means that your screen will display the picture of a cute animal, as illustrated in the screenshot below. You will not learn which situation was selected by the computer, neither before nor after your choice of option.

After making this choice, we will inform you about the number of participants in this session and ask you to guess the number of participants who have chosen "Recipient" and the number of participants who have chosen "Cute animal". Regardless of whether this period is selected or not for payment at the end of the session, you will earn 1 euro if your guess is correct, and 0 euro otherwise. Therefore, you should try to guess as accurately as possible. You will be informed on whether your guess is correct at the end of the session.

After you have chosen between "Recipient" and "Cute animal" and reported your guess, you will have to choose between "Option 1" and "Option 2". Your earlier choice between "Recipient" or 'Cute animal" determines the information you see on your screen before making your choice.

Summary of the Decisions

1. In Decision 1, you choose between Option 1 and Option 2.
2. In Decision 2, the program randomly selects which one of the two options cancels the donation.
3. You choose between the sets of information "Recipient" or "Cute animal".
4. You report your guess about the numbers of other participants in the session who chose 'Recipient" or "Cute animal".
5. Your screen displays the information you chose in step 3, and you choose between Option 1 and Option 2.

Please read again these instructions. If you have any questions, please raise your hand or press the red button. A comprehension questionnaire will be displayed on your screen.

OK

Part 2

In this part, you are randomly matched with 6 other participants to form a group of 7. There are two roles: Receivers and Senders. Receivers and Senders refer not to donations but to pictures and information, as explained below. All the participants in the group will first make decisions as Senders. Then, all of them will make a decision as Receivers. We first describe each role before explaining decision-making.

Choice of the Receiver

The Receiver has to choose between "Option 1" and "Option 2". The consequences from this decision are the same as in the second decision of part 1:

- "Option 1" results in €9 for the Receiver and "Option 2" results in €15 for the Receiver.
- The program picks randomly which one of the two options cancels the donation to GiveDirectly. Each option has 50 chances out of 100 to be picked.

The program randomly determines the consequences of each option independently for each Receiver. Thus, these consequences can differ across Receivers. Before making a choice between Option 1 and Option 2, the Receiver is not informed of the consequences of each option for GiveDirectly. However, s/he can obtain information from the Sender, as we now describe.

Choice of the Sender

With 80% chance, the Sender learns which one of the two options cancels the donation. With 20% chance the Sender does not learn the consequences of each option.

- If the Sender does not learn the consequences of each option for GiveDirectly, the program displays automatically the picture of a cute animal on the Receiver's screen before s/he makes his/her choice. The Receiver is not informed on the consequences of this option for GiveDirectly.
- If the Sender learns the consequences of each option for GiveDirectly, s/he has to choose between two types of information for the Receiver. If s/he chooses "Recipient", the Receiver will learn which one of the two options cancels the donation before choosing an option, and s/he will see the picture and the testimonial of a potential recipient of the donation. If the Sender chooses "Cute animal", the Receiver will see the picture of a cute animal, but not the consequences for GiveDirectly, neither before nor after the choice of option.

Decision Making in Part 2

In this part, all the participants first make two decisions in the role of a Sender in the case they are informed about the consequences of each option for GiveDirectly. Precisely, as a Sender, you have to choose between two types of information for the Receiver, either "Recipient" or 'Cute animal":

- in the case you learn that the donation to GiveDirectly is cancelled after "Option 1" , but not after "Option 2";
- and in the case you learn that the donation to GiveDirectly is cancelled after "Option 2", but not after "Option 1".

Then, all the participants will make a decision as Receivers. As a Receiver, you will have to choose between "Option 1" and "Option 2". Before you make your choice, the computer will randomly determine which option cancels the donation. It will also randomly pair you with a Sender in your group. The choice of the Sender between "Recipient" or 'Cute animal" determines the information you have about the consequences of each option.

Summary of the Decisions

1. You first decide as a Sender which picture and information to share if you are informed of the consequences of each option for Give Directly.
2. You are next a Receiver. You are randomly matched with a Sender.
3. You obtain the information chosen by this sender, "Recipient" or 'Cute animal".
4. You choose between Option 1 and Option 2.
5. You are paid based on your choice as a Receiver in case this part is selected for payment.

Please read again these instructions. If you have any questions, please raise your hand or press the red button.

Part 3

In this part, you are still matched with the same 6 other participants as in part 2. But now, participants are randomly assigned to one of the roles and will be identified with an ID. There are four Receivers and their IDs are R1, R2, R3, and R4. There are three Senders and their IDs are symbols (spade, diamond, club). We will communicate your role and your ID on your screen at the beginning of this part. This part has 25 identical periods and you will keep the same role and the same ID throughout this part. We now describe each of these periods.

Choice of the Receiver

In each period, the Receiver chooses between "Option 1" and "Option 2". The consequences from this decision are the same as before:

- "Option 1" results in €9 for the Receiver and "Option 2" results in €15 for the Receiver.
- The program picks randomly which one of the two options cancels the donation to GiveDirectly. Each option has 50% chance to be picked.

In each period, the program randomly determines the consequences of each option for GiveDirectly, independently for each Receiver. Thus, these consequences can differ across periods and across Receivers.

The Receiver is not informed about the consequences of each option for GiveDirectly. Before choosing between Option 1 and Option 2, s/he can receive information from the Sender.

Choice of the Sender

With 80% chance, the Sender learns the consequences of each option for GiveDirectly chosen by the program for each Receiver for the current period. If the Sender learns the consequences, s/he has to decide which set of information to share with the Receiver. As before, if s/he chooses "Recipient", i) the Receiver is informed which option cancels the donation before choosing an option, and ii) the picture of a potential recipient with his/her testimonial is displayed. If the Sender chooses "cute animal" the picture of a cute animal is displayed on the Receiver's screen and the Receiver does not learn which option cancels the donation, neither before nor after the

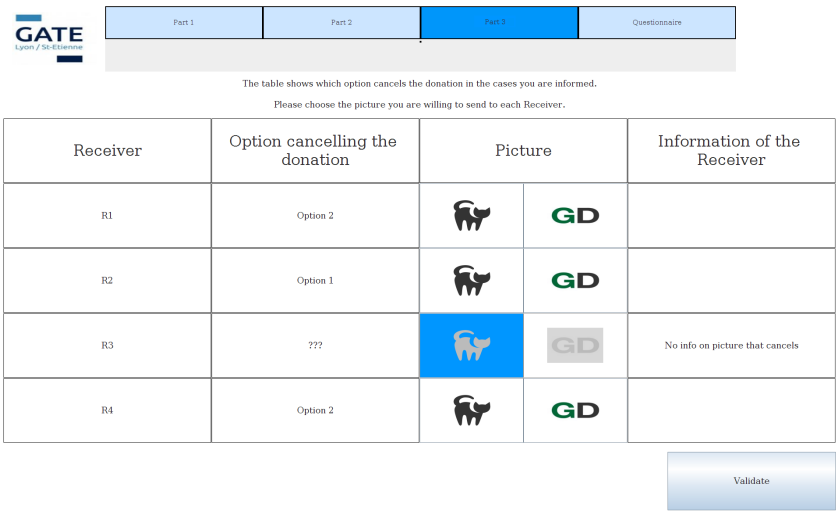

choice of option.

With 20% chance the Sender does not learn the consequences of each option and the picture of a cute animal is displayed automatically on the Receiver' s screen. The Receiver does not know whether the Sender has been informed or not.

The decision of the Sender is illustrated in the screenshot below. The first column of the table shows the ID of the Receivers (i.e. R1, R2, R3 or R4). The second column indicates for each Receiver, which option cancels the donation. The "???" sign indicates that the Sender did not receive information for participant R3; in this case, the sign of a cute animal is automatically pre-selected with no action from the Sender. In the next column, when informed, the Sender has to choose between "Recipient" and "Cute animal". In this example, the Sender makes three decisions, as s/he has information about the consequences of each option for three out of four Receivers. Once the Sender has made his/her decisions, the last column of the table indicates which information will be displayed on the screen of the Receiver.

As we explain now, the Sender can earn €10 for each Receiver that selects him/her in that period.

The Receiver Selects a Sender

Before the Receiver chooses an option, s/he has to select one of the Senders (spade, diamond, or club) for the current period. The selection decision is illustrated in the screenshot below, which reflects the screen the Receiver will see in the experiment. The example shows the decision in period 5. The first column shows the ID of each Sender. For each of the past periods, the screen shows which set of information each Sender shared with the Receiver (indicated by the symbol of an animal or the symbol GD for a recipient of GiveDirectly). The symbol of an animal reflects *either* that the Sender had no information, *or* that the Sender received the information and decided to share this picture; the Receiver cannot distinguish between these possibilities. The past choices of the Senders in the group are only visible to the Receivers, not to the other Senders. In this example, symbols have been chosen randomly.

When choosing a Sender, the Receiver only knows the past choices of the Senders but **not** those for the current period. The Receiver chooses a Sender by clicking on a box in the left column. The Receiver's information ("Recipient" or "Cute Animal") is determined by the selected sender's choice for the current period. The Senders are not informed of the choices of the Receivers.

Summary of the Decisions

1. For each Receiver, the program randomly selects which option cancels the donation.

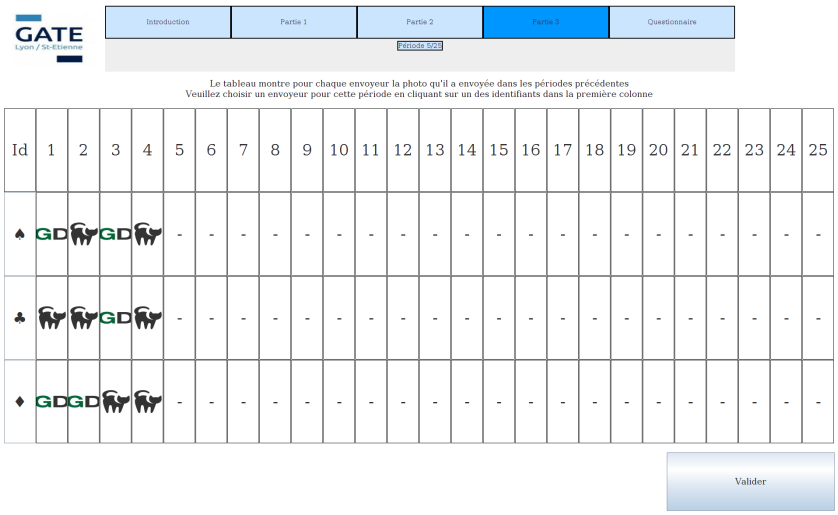

2. With 80% chance, the Sender learns the consequences of each option. If s/he learns the consequences, the Sender has to choose a set of information to share with the Receiver ("Recipient" or "Cute animal").
3. The Receiver sees information about sender's choices in earlier periods. S/he then chooses one of the Senders. The Receiver's information is determined by the choice of that sender ("Recipient" or "Cute animal") in the current period.
4. The Receiver decides between Option 1 and Option 2.
5. The same procedure applies for the 25 periods.

Summary of Earnings

The program randomly selects one of the 28 periods for payment (the two decisions in part 1, the decision as a Receiver in part 2 and the 25 periods in part 3). If a decision in part 1 or in part 2 is selected for payment, you will receive either €9 or €15, depending on whether you chose option Option 1 or Option 2 in that period. If a decision in part 3 is selected, payoffs for each player are as follows:

- **The Receiver** earns either €9 or €15, depending on the chosen option in that period.
- **The Sender** earns €10 for each Receiver that selected him/her in that round. Thus, the Sender minimally earns €0 if s/he has not been chosen by any Receiver, and maximally earns €40 if s/he has been chosen by the four Receivers.
- Finally, **GiveDirectly** will receive a donation of €15 for any choice in the selected period that does not cancel the donation.

End of the Session

At the end of part 3 a questionnaire will be displayed on your screen and then you will receive a feedback on your earnings in the session. On invitation of an experimenter, you will move into the payment room with your pre-filled receipt of payment and your computer tag.

Please read again these instructions. If you have any questions, please raise your hand or press the red button. A comprehension questionnaire will be displayed on your screen.

B Appendix: Additional Tables and Figures

Table B.1: Summary of Sessions

Session Number	Treatment	Nb Participants	Percentage of Females	Mean Age	Mean Payoff
1	ENDO	21	66.67%	29.09	18.76
2	ENDO	21	61.09%	23.38	18.52
3	ENDO	21	52.38%	23.67	18.81
4	ENDO	21	52.38%	23.43	19.33
5	ENDO	21	50.00%	21.45	18.33
6	ENDO	21	57.14%	20.29	18.76
7	ENDO	21	42.86%	21.47	19.05
8	ENDO	14	40.00%	20.87	18.21
9	EXO	21	42.86%	21.48	18.86
10	EXO	21	80.95%	20.76	17.67
11	EXO	21	57.14%	20.62	18.48
12	EXO	21	57.14%	23.52	15.48
13	EXO	21	61.90%	23.43	18.24
14	EXO	21	61.90%	21.68	18.90
15	EXO	14	57.14%	20.79	19.57
16	EXO	21	38.10%	23.10	19.14
Total	-	322	55.28%	22.50	18.49

Notes: The table reports the number of participants, the percentage of males, the mean age of the participants, and the mean participant's payoff in Euros, per session. The smaller number of participants in two sessions (one per treatment) is due to no show-up. The high mean age in session 1 is due to the presence of two participants aged 60 and 63.

Probability to choose the advisers — Multinomial models

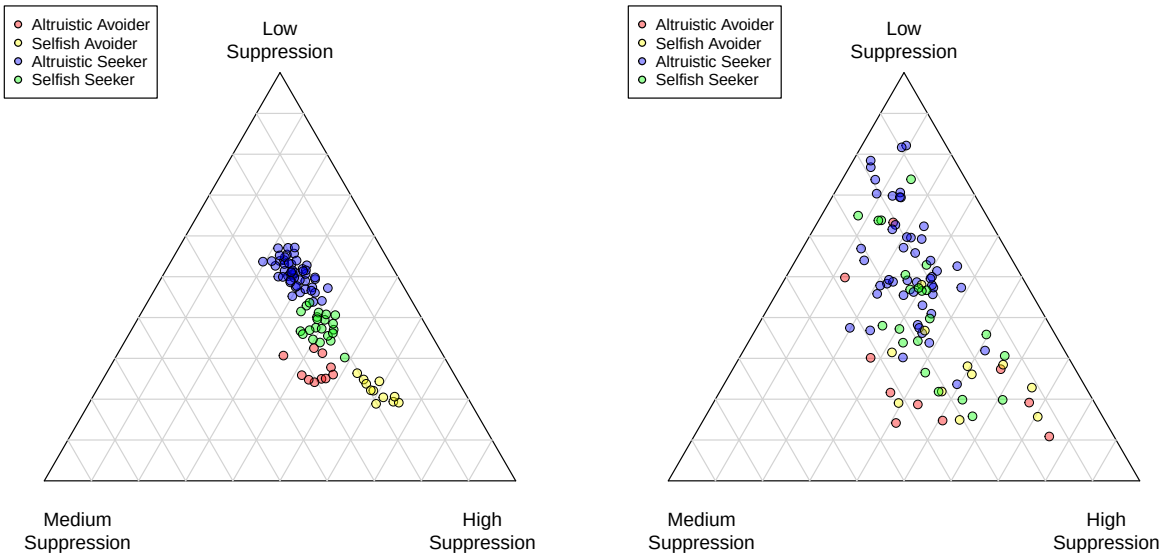

Figure B.1: Predicted probability of the decision makers' choice of adviser

Notes: The figures display the predicted probability to choose the adviser providing the highest, intermediate, and lowest level of suppression for each of the 92 decision makers. The figure on the left reports predictions based on Model 1 of Table 1 and the figure on the right reports predictions based on Model 2 of Table 1. The color of the dots captures the type of the decision maker elicited in Part 1.

Decision to switch adviser — Linear probability models

The decision to switch to another adviser conditional on the information received is informative of decision makers' information seeking or avoiding strategy. It provides an alternative way to look at information demand based on adviser history. Figure B.2 shows the fraction of decision makers that change adviser in part 3 in the ENDO treatment after adviser reported good, bad or no news. The left panel shows aggregate results, which demonstrates that switching rates are substantial and vary between 47% after bad news, 43% after good news, and 57% after no news. On aggregate, switching is highest after no news, in line with the idea that most people are information seekers.

The right panel shows a split by the "type" of decision makers, based on their revealed preferences in part 1. We use a two dimensional split. The first dimension opposes "Selfish" and "Altruistic" decision makers, based on the choice of these decision makers in the first decision with full information about the ethical trade-off. The second dimension splits decision makers into (information) "Seekers" or "Avoiders" depending on whether they chose to receive information or not in the individual decision under uncertainty. This right panel reveals that the aggregate results hide a lot of heterogeneity. In particular, the largest group (Altruistic - Seekers) clearly penalizes non-informative advice: the switching rate is 66% after receiving no news, 43% after bad news and 40% after after good news. By contrast, there is a smaller group (Selfish - Avoiders) that does the opposite: the switching rate is 27% after receiving no news, 53% after bad news and 44% after after good news. Moreover, except for the Altruistic - Seekers group, participants show a higher switching rate after bad than after good news, consistent with "shooting the messenger". This shows that decisions in the individual part of the experiment appear to have clear explanatory power for behavior in the interaction stage.

Table B.2 evaluates these results statistically in a linear probability model with the switching decision as an

independent binary variable and standard errors clustered at group level.[14]

In Model 1 the independent variables include dummies for receiving bad news or no news in the last round and interaction terms between these two variables and being classified as Selfish or as Avoider in part 1. Model 2 is the same as Model 1 but the regression includes a time trend and controls for a series of individual characteristics: age and grade obtained at the final high school exam (Baccalaureat) measured as a deviation to the average of all subjects in the experiment, a dummy for gender and the reported number of participations in past experiments. The results confirm that the baseline category (Altruistic - Seekers) is 26 percentage points more likely to switch adviser after no news, an effect that is almost entirely canceled in the group with Avoiders, which comprises 22 subjects (24% of all decision makers). We also see a significant effect for Selfish decision makers, who are significantly more likely to switch after bad news than altruistic subjects are.

Adding interactions between Selfish and Avoider to the previous models, Model 3 and Model 4 show that these results continue to hold. Decreased switching rates after no news are driven both by the Selfish and Altruistic Avoiders. For the Altruistic Avoiders, we can call this willful ignorance, as these subjects likely avoid bad news *in order* to avoid the cost of contributing to the charity. For the Selfish Avoiders, choosing uninformative advisers may not change their decision but may help them to avoid guilt from being confronted with the consequences of their actions.

Table B.2: Decision to switch advisers

	Model 1 Est. (S.E.)	Model 2 Est. (S.E.)	Model 3 Est. (S.E.)	Model 4 Est. (S.E.)
(Intercept)	0.409 (0.052)***	0.557 (0.062)***	0.405 (0.051)***	0.552 (0.061)***
GOOD news × Selfish	0.076 (0.091)	0.123 (0.082)	0.089 (0.097)	0.132 (0.084)
GOOD news × Avoider	-0.024 (0.103)	-0.023 (0.096)	0.005 (0.131)	-0.004 (0.128)
GOOD news × Selfish × Avoider	— —	— —	-0.062 (0.186)	-0.039 (0.188)
BAD news	0.025 (0.031)	0.026 (0.030)	0.023 (0.035)	0.020 (0.034)
BAD news × Selfish	0.105 (0.079)	0.164 (0.071)*	0.125 (0.098)	0.195 (0.080)*
BAD news × Avoider	0.025 (0.085)	0.025 (0.097)	0.064 (0.116)	0.094 (0.120)
BAD news × Selfish × Avoider	— —	— —	-0.084 (0.196)	-0.134 (0.195)
No news	0.268 (0.055)***	0.264 (0.060)***	0.256 (0.064)***	0.255 (0.070)***
No news × Selfish	-0.104 (0.076)	-0.034 (0.078)	-0.054 (0.102)	0.002 (0.107)
No news × Avoider	-0.249 (0.079)**	-0.252 (0.076)***	-0.169 (0.104)	-0.194 (0.088)*
No news × Selfish × Avoider	— —	— —	-0.170 (0.185)	-0.120 (0.179)
$Age - \overline{Age}$	— —	-0.000 (0.003)	— —	0.000 (0.003)
d(Male)	— —	-0.118 (0.052)*	— —	-0.118 (0.051)*
$BAC - \overline{BAC}$	— —	-0.000 (0.017)	— —	-0.000 (0.017)
Number of past participations in exp.	— —	-0.007 (0.017)	— —	-0.007 (0.017)
Period dummies	NO	YES	NO	YES
Number of observations	2208	2088	2208	2088
Number of clusters	23	23	23	23

Notes: These regressions are based on linear probability models. The binary dependent variable is the decision maker's choice to switch to another adviser in part 3 of the ENDO treatment. Robust standard errors clustered at group level are in parentheses. *d* for dummy variables. Period dummies are included with period 2 as the reference category). *** ≤ 0.001; ** ≤ 0.01; * ≤ 0.05.

[14] In some cases the coefficients do not precisely match the height of the bars in Figure B.2, as the former include demographic control variables, while the latter show pure frequencies.

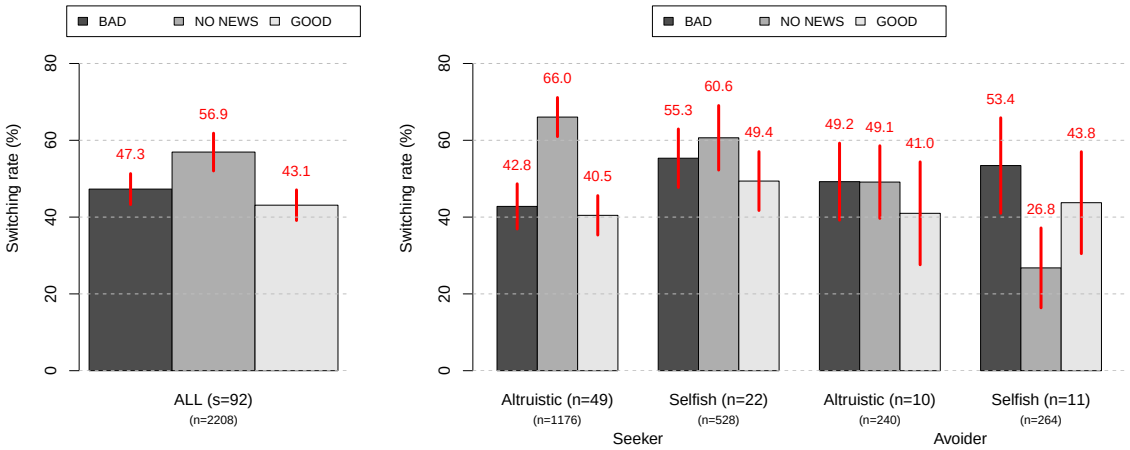

Figure B.2: Switching advisers

Notes: The figure displays the fraction of decision makers that change adviser after adviser reported good, bad or no news in the ENDO treatment. The left panel shows aggregate results. The right panel shows a split by decisions made by information seekers and information avoiders in part 1. Bars are standard errors based on the regression model 3 given in Table B.2.

C Illustration of Individual Demand and Supply Choices

Illustrations of the demand for information by decision makers with various d-statistics

The following three figures correspond to three decision makers with different d-statistics. In the figures, each line corresponds to one of the three advisers in the decision maker's group. The horizontal axis indicates the 25 periods in part 3 of the experiment. B is for bad news (option 2 cancels the donation); G for good news (option 1 cancels the donation); and an hyphen indicates that no news has been received. The colored items indicate which advisor has been selected in each period and which news has been revealed in the period after the adviser has been selected.

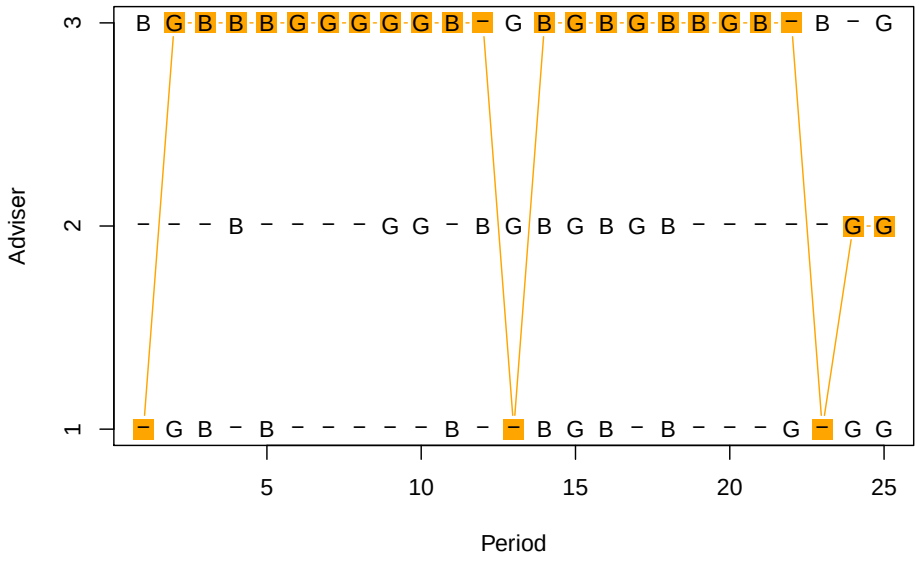

Figure C.1: Demand for information of a decision maker with d-statistic = 1.23

Note: Figure C.1 illustrates the case of an information seeker who sanctions any transmission of no news (d-statistic = 1.23).

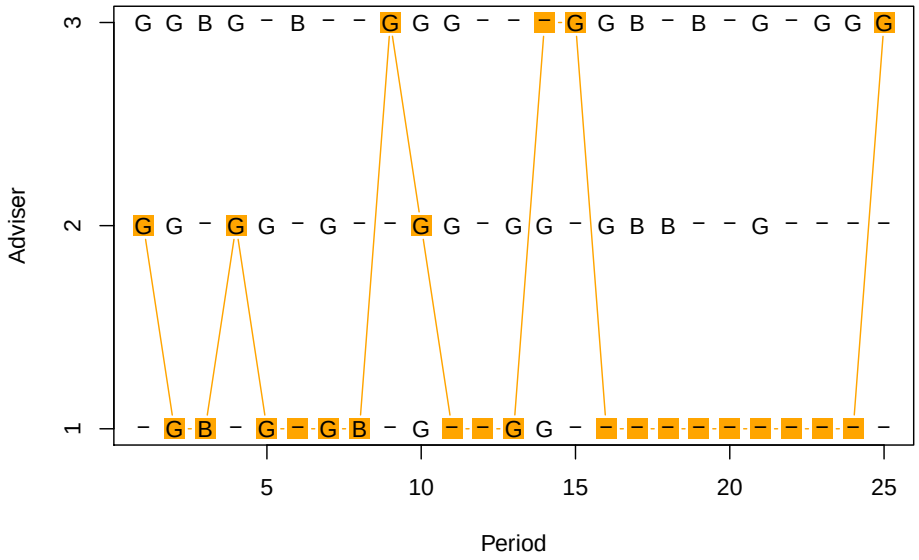

Figure C.2: Demand for information of a decision maker with d-statistic = 0.75

Notes: Figure C.2 illustrates also the case of an information avoider who is less able to establish a stable relationship with an adviser (d-statistic = 0.75).

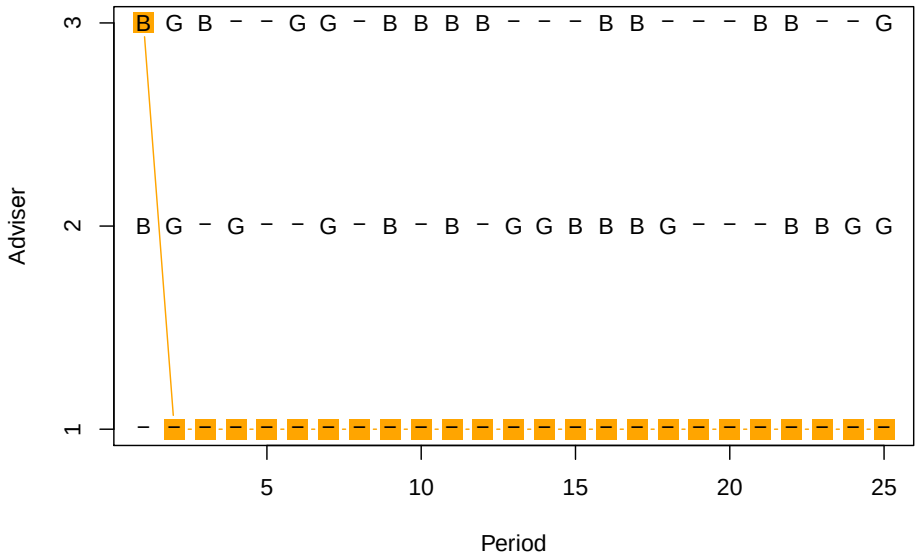

Figure C.3: Demand for information of a decision maker with d-statistic = 0.18

Notes: Figure C.3 illustrates the case of an information avoider who establishes a long term relationship with an adviser who never provides news (d-statistic = 0.18).

Illustrations of the supply of information by advisers with various s-statistics

The following three figures correspond to three advisers with different s-statistics. In the figures, each line corresponds to one of the four decision makers in the adviser's group. The horizontal axis indicates the 25 periods in part 3 of the experiment. B is for bad news (option 2 cancels the donation); G for good news (option 1 cancels the donation); and an hyphen indicates that no news has been sent (either because the adviser was not informed or because he or she decided not to send news). Colored letters indicate that the corresponding news has been sent and dark letters that the news has not been sent.

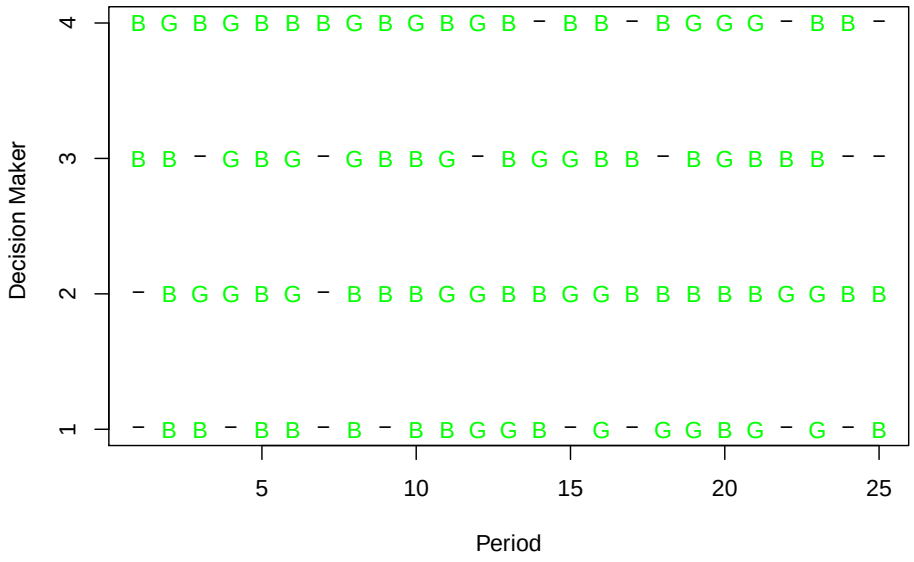

Figure C.4: Transmission of information to decision makers by an adviser with s-statistic = 0

Notes: Figure C.4 illustrates the case of an adviser that transmits all information (s-statistic = 0).

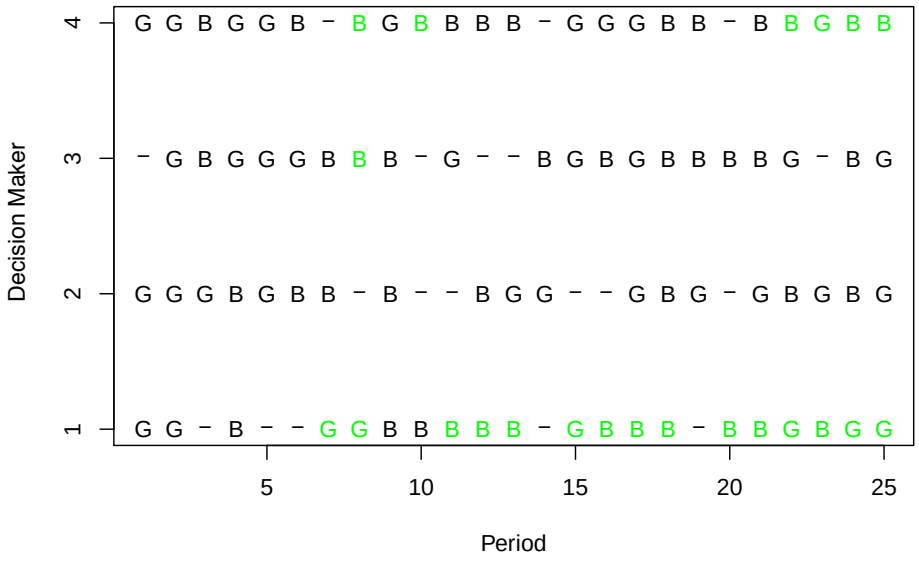

Figure C.5: Transmission of information to decision makers by an adviser with s-statistic = 0.66

Notes: Figure C.5 illustrates the case of an adviser that suppresses all types of news except for one decision maker (s-statistic = 0.66).

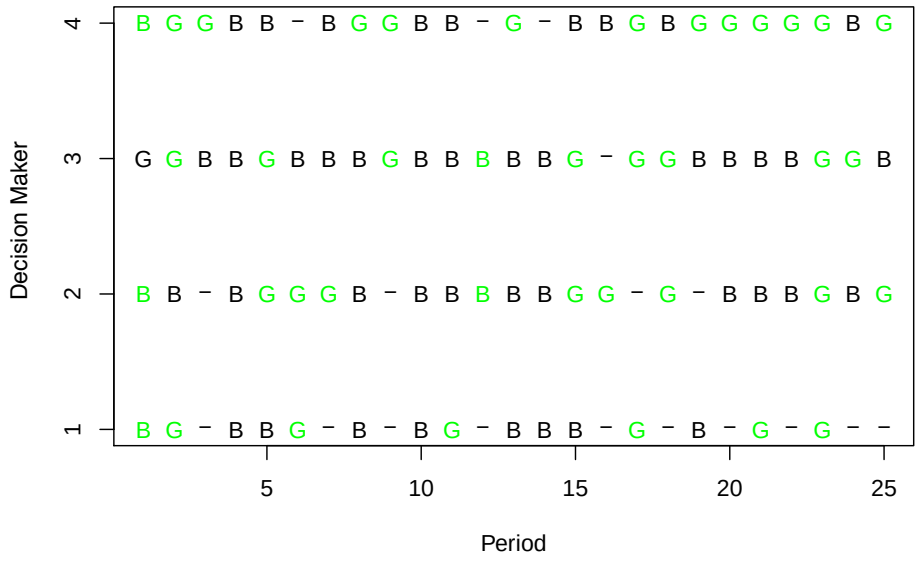

Figure C.6: Transmission of information to decision makers by an adviser with s-statistic = 0.89

Notes: Figure C.6 illustrates the case of an adviser that suppresses most bad news and does not discriminate among decision makers (s-statistic = 0.89). This adviser sends news almost only when they are good.