# The Social Construction of Ignorance: Experimental Evidence

Ivan Soraperra, Joël van der Weele, Marie Claire Villeval, Shaul Shalvi

# The Social Construction of Ignorance: Experimental Evidence*

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

July 17, 2022

**Abstract**

We experimentally study the social transmission of "inconvenient" information about the externalities generated by one's own decision. In the laboratory, we pair uninformed decision makers with informed senders. Compared to a setting where subjects can choose their information directly, we find that social interactions increase selfish decisions. On the supply side, senders suppress almost 30 percent of "inconvenient" information, driven by their own preferences for information and their beliefs about the decision maker's preferences. On the demand side, about one-third of decision makers avoids senders who transmit inconvenient information ("shooting the messenger"), which leads to assortative matching between information-suppressing senders and information-avoiding decision makers. Having more control over information generates opposing effects on behavior: selfish decision makers remain ignorant more often and donate less, while altruistic decision makers seek out informative senders and give more. We discuss applications to information sharing in social networks and to organizational design.

**Keywords:** Social interactions, information avoidance, assortative matching, ethical behavior, experiment

*JEL:* C91, D82, D83, D91

*An early version of this paper circulated under the title "Shooting the Messenger? Supply and Demand in Markets for Willful Ignorance." The authors thank the editor and three anonymous referees for their constructive and helpful comments. They are grateful to participants at the Beliefs workshop in Munich, the IMEBESS Conference in Utrecht, the ESA conference in Dijon, and the FELIS CORTEX workshop in Lyon for useful comments, and thank Björn 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). The research was further supported by a Dutch Science Association grant (NWO VI.Vidi.195.137) and the European Research Council (ERC) grants (ERC-StG-637915 and ERC-CoG-865931).

†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.

¶CREED, Univerity of Amsterdam. Email: S.Shalvi@uva.nl.

# 1 Introduction

Many of our decisions impact others, sometimes in a visible way, more often in invisible ways. Such invisible or ambiguous impact occurs notably when our consumption decisions contribute to environmental damage or suffering by humans and animals further down the supply chain. Informing ourselves about these impacts is crucial for making pro-social decisions. However, such information is potentially "inconvenient", as it may highlight trade-offs between personal profits and moral behavior. A growing literature on willful ignorance shows that people are sometimes reluctant to acquire ethical information, and may actively avoid it to excuse their selfish behavior.[1]

While the literature has focused on individual decisions, an unexplored issue is how inconvenient information is shared and consumed in social interactions. This is an important question, as everyone regularly supplies information to friends, colleagues and contacts on social media, and must sometimes decide whether to truthfully relay inconvenient information. Different motives may play a role in the decision to share such information, such as paternalism and a feeling of duty to tell the truth. Other motives might be more strategic, as people may cater to (perceived) demand for information or willful ignorance. Thus, supply decisions may depend on whether others reward provision of truthful information, or prefer to seek out those who shield them from inconvenient truths.

In this paper, we investigate information sharing and its effect on ethical decision making. To do so, we conduct a laboratory experiment in the context of ethical dilemmas with uncertainty. Participants in the role of "decision makers" face a choice between actions that differ in their profitability for the decision maker and carry a risk of entailing a negative externality for a charity. Each decision maker is matched to another subject in the role of "sender", who is likely to have superior information about the consequences of each decision for the charity. The sender can choose to disclose this information to the decision maker, or send an irrelevant distraction (a picture of a cute animal) instead. In this setting, we study the willingness of senders to supply inconvenient information that highlights a trade-off between profits and externalities, and how this affects subsequent decision making.

[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 (Larson and Capra, 2009; Nyborg, 2011; Conrads and Irlenbusch, 2013; Grossman, 2014; Feiler, 2014; Bartling et al., 2014; Kajackaite, 2015; van der Weele, 2013; Grossman and van der Weele, 2017; Espinosa and Stoop, 2019; Serra-Garcia and Szech, 2019). See Vu et al. (2022) for a meta-analysis. For an analysis of the neural correlates of information seeking or avoidance, see Charpentier et al. (2018); Sharot and Sunstein (2020). Related work shows how self-serving interpretations of risk and ambiguity increase selfishness in sharing decisions (Haisley and Weber, 2010; Di Tella et al., 2015; Exley, 2015; Garcia et al., 2020). Freddi (2019) provides evidence of information avoidance from the field. There is also ongoing and inconclusive research about willful ignorance in product markets, with Bartling et al. (2015) finding little evidence, whereas Momsen and Ohndorf (2020) and particularly Ehrich and Irwin (2005) and Momsen and Ohndorf (2019) find more positive evidence.

To study the strategic aspects of ethical information sharing, we vary the possibility for decision makers to *choose* their information sources. Across different treatments, we implement either random matching between informed senders and decision makers, or, like in a social network, give decision makers the possibility to choose senders on the basis of their past information sharing decisions. To mimic senders' incentives to attract clients or followers on social media, senders earn money from each matched decision maker in each treatment. This treatment allows us to study whether decision makers "shoot the messenger" of bad news, and whether this affects the supply of inconvenient information.

We find evidence that ignorance results from both the demand and the supply side. Despite the salience of the ethical dilemma and the presence of relevant information, almost 30 percent of shared information consists of irrelevant animal pictures. A majority of senders suppresses inconvenient information sometimes, and a quarter does this most of the time. Suppression correlates with senders' own preferences and attitudes toward information, suggesting paternalistic motives. The strategic motive to attract more decision-makers appears less important in our setting. We do also find some evidence for a wish to entertain the decision maker: in a condition where we eliminate cute animal pictures as a distraction, we see a reduction in suppressing of information about the recipient in cases where there is no ethical trade-off.

We also find a clear demand for ignorance. If decision makers have the choice to do so, about one-third "shoots the messenger", and shifts toward senders who suppress inconvenient information. These patterns of demand and supply lead to assortative matching: Compared to random matching, decision makers who can choose their senders are more likely to match with senders with similar preferences for information in the social dilemma, as these are more likely to supply their preferred information content.

Following the previous literature on willful ignorance, we also implement a condition where we eliminate the sender, so decision makers directly choose their own information. We find that selfishness goes down in this setting compared to social interactions. The main reason is that altruistic decision makers typically choose to be informed when they can do so themselves, but do not always obtain the necessary information to make an ethical decision when they interact with senders. More generally, giving decision makers more control over their information, by letting them choose it directly or because they can pick their sender, has ambiguous effects on decision making. It leads altruistic senders to obtain more information and become more prosocial. By contrast, selfish senders choose less informative senders and become slightly less prosocial. On aggregate, these two effects partially offset each other.

Thus, we show that social interactions facilitate the production of willful ignorance, result-

ing from a complex set of motives. Our study is explorative and does not test a particular model, but the results speak to various applications where people can choose to share ethical information. The fact that irrelevant distractions make up a substantial share of information content is in line with casual observations on social media platforms, and may help explain persistent misperceptions about uncomfortable topics, like climate change, children's labor, or animal suffering. Our results also provide confirmatory evidence for the segmentation of information into "filter bubbles" (e.g., Aiello et al., 2012), where users match depending on the type of information that they like to see. The suppression of ethically relevant information is a novel finding that deserves follow-up investigation, given the ubiquitousness of sharing decisions both on and offline.

A second application relates to organizational design. Effective decision making in organizations requires that advisers and consultants give executive decision makers relevant and unbiased information, rather than being "yes-men". If we allow some extrapolation, our matching treatments can be viewed as comparing a fixed bureaucracy, or "deep state", with a system where executives bring in their own advisers. We find that both systems lead to similar amounts of prosocial behavior. However, the matching scheme determines *who* has more power over the decision. Under random matching, the power of the executive is reduced, as we find that the decision makers' preferences in the social dilemma are less predictive of their decisions in this case.

Our paper makes several contributions to the existing literature. First, we contribute to the literature on individual information avoidance, by studying not just demand, but also the supply of information. Previous studies have considered advice in social dilemmas, such as Schram and Charness (2015) and Coffman and Gotthard Real (2019), where unlike in our study, advisers do not have an informational advantage and can only express their opinion.[2] 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. Several papers investigate the role of image concerns in the supply of information about externalities (Foerster and Van der Weele, 2018; Foerster and Van der Weele, 2021; Bénabou et al., 2018). As far as we know, we are the first to study the matching between the supply and demand for ethical information.[3]

Second, we relate to a growing literature on group decisions and the dilution of responsibil-

[2] A number of studies have looked at settings where advisers have an informational advantage, but these studies typically focus on conflicts of interest between advisers and decision makers, looking for instance at the disclosure of such interests (e.g. Ismayilov and Potters, 2013) or cognitive dissonance of advisers (Chen and Gesche, 2016; Gneezy et al., 2020).

[3] 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 as a strategy of ingratiation for agents who 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.

ity. While the contributions are numerous, prominent examples include Dana et al. (2007) who document how pairs of subjects are more selfish than individuals, and Falk and Szech (2013) who 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. 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. More generally, Charness and Sutter (2012) provide survey evidence that groups make more selfish decisions than individuals. We contribute to this literature by showing how informed and uninformed players collaborate on information suppression. Our results show that social interactions do not mitigate willful ignorance, and may make it worse.

The remainder of this paper is organized as follows. Section 2 introduces our experimental design and procedures and presents our main behavioral conjectures. Section 4 presents and discusses our main results. Section 5 discusses an extension of our design where we suppress cute animal pictures. Section 6 concludes.

# 2 Design

The experimental design consists of three parts and two treatments. The first two parts are identical for all participants, and serve to elicit some characteristics of interest from each participant and familiarize them with the experimental setting. The third and main part differs across treatments. We describe each part in turn. The instructions are available in Appendix A.

## 2.1 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 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 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

[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.

instructions in Appendix A).

After this decision, 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. The diagram in Figure 1 provides a summary of the payoffs in the two options. Participants are not informed of the outcome of the random draw. However, before making their choice, they have to choose whether they want to be 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. Note that 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's and the charity's interests.

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".

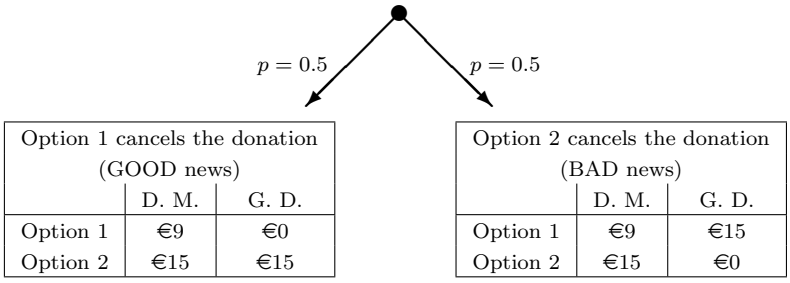


*Figure 1: Summary of the decision maker's and charity's payoffs*

*Notes*: The diagram summarizes the payoffs of Option 1 and Option 2 for the decision maker (D.M.) and the charity (G.D.) when Option 1 cancels the donation (GOOD news) and when Option 2 cancels the donation (BAD news). Note that the notions of Good and Bad news were not used in the experiment.

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. As we explain below, this elicitation serves to better understand the strategic motives of the senders in the experiment.

## 2.2 The role of senders

In the next parts of the experiment, decision makers face a decision that is identical to that 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. However, decision makers are not informed of the outcomes of this draw and cannot choose information about the consequences of their decision directly. Instead, decision makers' information provision depends on the decisions of other subjects in the role of "senders", who may have information about the consequences of the decision that can be transferred to decision makers. Subjects are randomly matched in groups of seven. Within each group, three participants are randomly assigned the role of sender (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, and referred to as "receivers"). Subjects keep the same role and identifier throughout the experiment, which consists of 25 rounds, each with an identical structure.

In each round, information is transferred from a sender to a decision maker if a) the sender decides to share information about the consequences of the decision maker's action and b) there is a match between sender and decision maker within the group. The matching protocol is either random or by choice, depending on the treatment conditions that we discuss below. To maximize the amount of data about sender's sharing behavior, senders decided whether to share information with each of the four decision makers in their group, before they knew whether they would be matched with that decision maker (see screenshots in Appendix A). These conditional sharing decisions are implemented only for those decision makers that match with the sender. If the sender decided to share, the potential beneficiary's picture and testimonial and information on which option cancels the donation are sent to the matched decision maker. If the sender decided not to share, a picture of a cute animal is sent to the decision maker without information on which option cancels the donation, neither before nor after his or her choice of option.

We also implemented an important caveat to the information transfer. In 20 percent of cases (drawn independently for each decision maker in each round), senders were not informed about the consequences for the decision maker. In this case, the picture of a cute animal is automatically sent to the decision maker when a match is established, without a choice from the sender. This design feature served two purposes. First, such wiggle room is a realistic feature of many situations, as sources or advisers are unlikely to always know the truth. Second, it allows senders some wiggle room, as non-disclosure is not exclusively explained by an unwillingness to share information. Thus, if a sender and decision maker coordinate on the suppression of information between them, they can "pretend" that this ignorance is externally imposed. Accordingly, when a decision maker can see the picture of a cute animal on the screen, he or she does not know whether the sender selected it or whether he or she was uninformed.

## 2.3 Part 2: Role familiarization tutorial

The main interaction between senders and decision makers will be repeated multiple times, and the resulting decisions will be the focus of our analysis. However, because the structure is somewhat complex, we first included a short tutorial. This tutorial, which we refer to as Part 2, served to practice with the roles of both sender and decision maker. It also helps subjects understand the potential motives of participants in both roles, regardless of the role subjects will eventually play in the main part of the experiment.

In Part 2, all participants first play in the role of a sender. Participants 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. We used the strategy method as we wanted them to practice both decisions. After making their choices as a sender, all participants play in the role of a decision maker. Each participant is randomly matched with another player in the group of seven. The decision maker's information depends on the choice of this other participant when he or she played in the role of a sender. Depending on the decision of this sender, the decision maker screen either displays the picture of a cute animal, or indicates which option cancels the donation and displays the picture and testimonial of a potential recipient before the decision maker's choice of option.

To make subjects take the exercise seriously, we allowed this part to be selected for payment. If this happened, players were paid based on their choice of option as a decision maker (see Section 2.6). We did not incentivize sender behavior in this part. The main sender incentives we will consider are the selection choices of decision makers, but this was not an option in the tutorial, there senders and decision makers were matched randomly.

## 2.4 Part 3: Main experiment and matching treatments

In the third and main part of the experiment, participants remain matched with the same six other participants as in Part 2 and make choices in each of 25 periods. Subjects are randomly allocated to one of two treatment conditions, RANDOM and CHOICE, which differ only in the matching process of senders and decision makers.

In the RANDOM treatment, before choosing an option the decision maker is randomly matched by the program with one of the senders for the current period. He or she receives the information shared by this sender for the current period and chooses one of the two options. In the CHOICE treatment, each decision maker has to select one of the senders before choosing an option. Thus, in the CHOICE treatment subjects can select which type of sender they prefer, either those who are likely to share information or those who may help them remain willingly ignorant.

To facilitate the choice of a sender, the decision makers get some information at the begin-

ning of each round about the senders' behavior in previous rounds. A history box displays a symbol for each type of information sent to him or her by each sender in each of the *previous* periods. Symbols are either 'GD' for GiveDirectly—when the sender sent information with the picture and testimonial of a potential recipient—or the symbol of an animal—if the sender had no information or he or she received the information and decided to send the picture of the cute animal (see screenshots in A). The past choices of the sender in the group are only visible to the decision makers, not to the other senders. The history box was provided in both the CHOICE and RANDOM treatment, although in the latter the information could not be used to select a sender.

Senders in both treatments are paid €10 for each decision maker they are matched with, either exogenously in the RANDOM treatment or endogenously in the CHOICE treatment. This gives senders in the CHOICE treatment an incentive to attract decision makers. To facilitate this, senders learn the identifier of the decision makers that were matched with them in that period. Senders can use this information to determine their strategy in subsequent rounds. For instance, they can continue to send the same information to subjects who matched with them, and experiment with sending different types of information to the others in order to attract them. To reduce the complexity of the environment, senders are not informed about the option eventually chosen by the decision makers, or the information shared by other senders in the group. This obviously limits their ability to learn about their strategic environment, so in future research it would be interesting to study how senders react to such information.

### 2.5 Follow-up treatments

After gathering the data on the RANDOM and CHOICE treatment and writing an initial working paper, we conducted two follow-up experiments.[5] Following the existing literature on willful ignorance, an INDIVIDUAL treatment investigates what happens if subjects are in control of their own information choices instead of relying on a sender. This treatment serves as a control for the presence of social interactions in both CHOICE and RANDOM. In INDIVIDUAL, the interaction in Part 3 is replaced by an individual decision. This decision thus repeats 25 times the decision in Part 1 under uncertainty. These 25 choices only differ in the realization of the payoff structure, the sample testimonial of the recipient and the cute animal pictured, where the latter differs in each round. As in the other treatments, there is a 20 percent chance in each round that recipient information is not available, even if the decision-maker chose to reveal it. Since there was no sender, we dropped Part 2 from the experiment.

In addition, we conducted follow-up treatments to understand the importance of cute animal pictures, which we report in Section 5. These NoCAP treatments are a replication of the RANDOM and CHOICE treatments discussed above without the cute animal picture when

[5] We thank two anonymous referees for suggesting these treatments.

information is not provided.

## 2.6 Procedures

All sessions were conducted at GATE-Lab, Lyon, France. For our original treatments, we ran 16 sessions (8 for the RANDOM treatment and 8 for the CHOICE treatment). The 322 participants (161 in the RANDOM treatment and 161 in the CHOICE treatment) are mainly students recruited from the local engineering, business and medical schools, using Hroot (Bock et al., 2014). Our main treatments (CHOICE and RANDOM) were conducted in the Summer and Fall of 2018 (June to November). Regarding the follow-up treatments, we ran 5 sessions of the INDIVIDUAL treatment with a total of 75 participants in November and December 2021.[6] For the treatments without the cute animal picture (NoCAP) (see details in Section 5), we ran 16 sessions in November and December 2021, with 154 participants in the RANDOM treatment and 161 participants in the CHOICE treatment.[7] In total, the experiment involved 712 subjects. Table B.1 in Appendix B gives a summary of the sessions and the fraction of females and mean age in each session. The experiment was developed in Java.

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 sender earned €10 for each decision maker he or she was matched with in that period (thus, the sender 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 all 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 €19.04 (standard deviation, S.D. hereafter, = 6.25), including a €5 show-up fee. Payments for the main experiment were made in cash, in a separate room and in private;

[6] Due to a coding error, we did not record behavior in two sessions, which were rerun in February 2022.

[7] In one session of the RANDOM treatment we are one market (7 participants) short of the original numbers, due to insufficient show-up of participants.

because of the COVID-19 pandemic, payments for the follow-up experiment had to be made by bank transfers.

## 3 Behavioral conjectures

We designed our experiment to study the drivers of supply and demand for ethically relevant information and their matching in a social interaction, as we explain in more detail below. We consider this an initial exploration, as it is hard to capture our setting in a formal model for two reasons. First, there is a multitude of psychological motives involved in both demand and supply of ignorance, relating to self-interest, social preferences, self-image, paternalism, etc. Second, the dynamic nature of our interactions allows for a multitude of strategies. Nevertheless, we form a number of conjectures about the behavior we expect to observe in Part 3.[8]

The behavior of subjects will depend on their preferences and their attitudes toward information. Thus, we distinguish between different types of subjects depending on their motivations. First, subjects may differ in their preferences over the payoffs to themselves and the charity. *Selfish* subjects are only motivated by the maximization of their individual payoff, whereas *Altruistic* decision makers are willing to make sacrifices on behalf of the charity. A measure of these preferences is provided by the first decision in Part 1 of the experiment, where subjects make individual decisions with full information. Second, there are further psychological motivations like self-image and guilt that may affect information attitudes. Previous literature has shown that roughly one third of subjects can be classified as "reluctant" altruists (Dana et al., 2006, 2007; Lazear et al., 2012). These agents want to choose the selfish option, but also want to maintain a positive self-image or avoid cognitive dissonance from being explicitly selfish. Remaining uninformed may serve as an excuse and help maintain self-image while also reaping a profit (Grossman and van der Weele, 2017). The second decision in Part 1 of the experiment shows whether people prefer to avoid information, and allows us to classify subjects as information *Avoiders* or *Seekers*.

How do such heterogeneous motives affect decisions in the experiment? We first consider information supply. Altruistic senders who care about the charity would always want to disclose bad news in order to inform the decision maker of the possible trade-off. By contrast, Selfish senders may deviate from full disclosure for several reasons. Given that senders are paid for each matched decision maker, there may be a strategic motive. If senders anticipate a sufficient demand for ignorance in the CHOICE treatment, they may suppress bad news in order to attract more decision makers. Thus, we expect that in the CHOICE treatment, senders'

[8]Note that we have originally pre-registered an early version of the project. The project, however, has significantly developed. The analysis deviated from our original plan and we did not preregister the follow-up experiments. Accordingly, we no longer consider part of our conjectures pre-registered. For transparency purposes, readers can find the pre-registration document here: https://aspredicted.org/blind.php?x=dr7743.

suppression will be correlated with their beliefs about the decision makers' preferences for information. This effect should be eliminated in the RANDOM treatment. However, senders may also want to do the decision makers a "favor" by not confronting them with a difficult ethical decision. In this case, we may see a (weaker) correlation also in the RANDOM treatment.

**Conjecture 1.** *(Supply side). Selfish senders will suppress more inconvenient information if they believe more decision makers want to be uninformed. This effect is stronger in the CHOICE than in the RANDOM treatment.*

There are additional motives for suppression that lead to further hypotheses. For instance, if the sender wants to impose his or her own information preference on the decision maker (paternalism), one would expect that Avoiders are more likely to suppress news. A sender may also aim to entertain the decision maker with pictures of cute animals.

We now turn to information demand by decision makers. Selfish decision makers should choose Option 2 in all treatments. Because information does not change their decision, standard economic theory would predict that Selfish participants are indifferent between information sources. However, Selfish Avoiders, following the guilt-avoidance logic outlined above, are motivated to select uninformative senders that sent animal pictures in previous periods. By contrast, Altruistic decision makers should choose Option 2 if and only if they are informed that Option 1 cancels the donation, in order to avoid the risk of cancelling the donation by their decision. Indeed, having information about the consequences is necessary to be altruistic, since in the absence of information, either option is equally likely to cancel the donation. Thus, in the CHOICE treatment Altruistic decision makers should select informative senders, *i.e.*, those who in the past periods were most likely to disclose information about the charity. By contrast, Altruistic Avoiders, who would like an excuse in order to behave selfishly, should select uninformative senders. This analysis leads to the following behavioral conjecture.

**Conjecture 2.** *(Demand side). There is a demand for ignorance. Decision makers who avoided information in Part 1 ("Avoiders") will seek out senders with uninformative messages.*

We also make a conjecture about differences in matching patterns across treatments. If Avoider (or Selfish) type senders are more likely not to disclose information, and Avoider (or Selfish) type decision makers are more likely not to look for information, we should see assortative matching of types, even if preferences are not directly observable to the other side.

**Conjecture 3.** *(Assortative matching). There is assortative matching in the CHOICE treatment, with Avoider decision makers more likely to match with Avoider senders.*

When it comes to ethical behavior, it is hard to make general predictions, as several mechanisms are at play. In particular, our treatments may have different and opposing effects on the different types of decision makers. Avoider type individuals may choose senders who suppress information in the CHOICE treatment and hence obtain less information than in the RANDOM treatment, and yet less in the INDIVIDUAL treatment since in this treatment they can remain ignorant for sure. The reverse is true for Seeker types. Because of these divergences in information acquisition, we expect that Altruistic and Selfish types will also diverge in the type of behavior. In particular, since information is necessary to behave altruistically, Altruists will become more altruistic when they have more control over their information, whereas selfish people will become even more selfish, since they are less often confronted with information about recipients and hence feel less pressure to become prosocial.

**Conjecture 4.** *(Heterogeneity). When control over information choices increases from RANDOM to CHOICE to INDIVIDUAL, we expect that*

1. *Preferences for information (Avoider vs. Seeker) will have a stronger impact on the consumption of ignorance.*
2. *Preferences for prosocial behavior (Selfish vs. Altruist) will have a stronger impact on ethical behavior.*

# 4 Results

We first give an overview of the type of information transmitted in our main treatments and its impact on decisions. We then turn to analyze the supply and demand of information in the CHOICE and RANDOM treatments, as well as the matching of different types of decision makers and senders in the CHOICE treatment and the impact on their behavior. Throughout, we will use "good news" to refer to messages that show no ethical trade-off (Option 1 cancels the donation), "bad news" for messages that show such a trade-off (Option 2 cancels the donation), and "no news" to uninformative animal pictures.

Table 1 shows an overview of the distribution of the different types (Seekers vs. Avoiders and Altruists vs. Selfish), as measured by the subjects' behavior in Part 1. Unfortunately, these distributions are not balanced across treatments ($\chi^2(12) = 35.58$, $p < 0.001$), so when we compare aggregate outcomes across treatments, we will control for the types of decision makers.

*Table 1: Distribution of types by treatment*

| Type | (CAP) Cute animal picture | | | (NoCAP) No cute animal picture | |
|---|---|---|---|---|---|
| | RANDOM | CHOICE | INDIVIDUAL | RANDOM | CHOICE |
| Selfish - Avoider | 11 ( 6.8%) | 18 (11.2%) | 10 (13.3%) | 7 ( 4.5%) | 19 (11.8%) |
| Selfish - Seeker | 35 (21.7%) | 42 (26.1%) | 25 (33.3%) | 50 (32.5%) | 70 (43.5%) |
| Altruistic - Avoider | 18 (11.2%) | 16 ( 9.9%) | 3 ( 4.0%) | 15 ( 9.7%) | 7 ( 4.3%) |
| Altruistic - Seeker | 97 (60.2%) | 85 (52.8%) | 37 (49.3%) | 82 (53.2%) | 65 (40.4%) |
| Total | 161 (100.0%) | 161 (100.0%) | 75 (100.0%) | 154 (100.0%) | 161 (100.0%) |

*Notes*: The table reports the distribution of preference types (in rows) by treatment (in columns). Types are defined based on their choices in Part 1: Selfish (Altruistic) chose Option 2 (Option 1) in decision 1 of Part 1, i.e., when informed that the state of the world is bad; Avoider (Seeker) chose to obtain (avoid) information in decision 2 of Part 1. Columns 2, 3, and 4 report the distribution in the treatments of the experiment where cute cute animal pictures were present (CAP); Columns 5 and 6 in those where we eliminated the cute animal pictures (NoCAP). Number of subjects with percentages in parentheses.

## 4.1 Ethical behavior and information consumption

As an overall measure of ethical behavior, we consider decisions in the "bad" state, that is, where there was an ethical trade-off (regardless of whether decision makers were informed or uninformed). Overall, we find very similar fractions of 62.4%, 62.2% and 61.0% of selfish decisions in the CHOICE, RANDOM and INDIVIDUAL treatment, respectively. However, since the distribution of subjects' preferences is not well balanced (see Table 1), we should control for subjects' preferences. Column 1 of Table 2 shows that once we control for this difference, we find that selfish behavior is significantly lower in the INDIVIDUAL treatment by about 9 percentage points, while the difference between RANDOM and CHOICE is not significant (Wald test: $\beta = 0.083 - 0.102 = -0.019$, $p = 0.688$). Column 2 of Table 2 shows that this result hides a lot of heterogeneity between different types of decision makers, which we discuss in more detail in Section 4.5.

To understand the origins of ethical behavior, we look at news consumption, as well as selfish behavior conditional on news consumption. Since information was not always available, the expected distribution of information available to senders is 40% good news, 40% bad news and 20% no news.[9] Thus, if senders transmitted all information or if decision makers selected only senders who did so, this should be the distribution of information consumed by the decision makers in Part 3. The left panel of Figure 2 shows that the actual distribution of information observed by the decision makers differs starkly from this benchmark ($\chi^2(2) = 594.16$, $\chi^2(2) = 668.54$, $\chi^2(2) = 265.52$ in the CHOICE, RANDOM, and INDIVIDUAL treatments, respectively).[10] With a prevalence of 40.3% in the CHOICE treatment and 41.6% in the RANDOM treatment, no news is most commonly consumed, whereas good and bad news are both

[9]The realized frequencies are: 38.3%, 40.7%, and 21.0% in the CHOICE treatment; 39.7%, 40.8%, and 19.5% in the RANDOM treatment; and 38.2%, 41.0%, and 20.8% in the INDIVIDUAL treatment.

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

*Table 2: Unethical outcomes and ignorance consumption*

| | Unethical outcomes | | Ignorance consumption | |
|---|---|---|---|---|
| | Model 1 | Model 2 | Model 3 | Model 4 |
| | Est. (S.E.) | Est. (S.E.) | Est. (S.E.) | Est. (S.E.) |
| (Intercept) | 0.846 (0.056)*** | 0.916 (0.053)*** | 0.311 (0.047)*** | 0.307 (0.056)*** |
| d(CHOICE) | 0.083 (0.043)° | -0.046 (0.053) | 0.052 (0.039) | 0.031 (0.064) |
| d(RANDOM) | 0.102 (0.040)* | -0.037 (0.043) | 0.069 (0.044) | 0.075 (0.069) |
| Avoider DM | 0.087 (0.038)* | 0.055 (0.047) | 0.145 (0.037)*** | 0.350 (0.080)*** |
| Avoider DM × d(CHOICE) | | 0.072 (0.089) | | -0.231 (0.099)* |
| Avoider DM × d(RANDOM) | | -0.012 (0.075) | | -0.329 (0.091)*** |
| Altruistic DM | -0.486 (0.032)*** | -0.633 (0.046)*** | -0.064 (0.027)* | -0.121 (0.054)* |
| Altruistic DM × d(CHOICE) | | 0.206 (0.069)** | | 0.109 (0.070) |
| Altruistic DM × d(RANDOM) | | 0.246 (0.064)*** | | 0.095 (0.071) |
| $Age - \overline{Age}$ | -0.005 (0.002) | -0.005 (0.002)* | -0.004 (0.002)* | -0.004 (0.001)** |
| *d(Male)* | -0.001 (0.033) | 0.006 (0.031) | -0.007 (0.026) | -0.008 (0.025) |
| $BAC - \overline{BAC}$ | 0.005 (0.008) | 0.004 (0.008) | -0.004 (0.007) | -0.004 (0.007) |
| *# of past participations in exp.* | 0.008 (0.011) | 0.012 (0.010) | 0.006 (0.009) | 0.006 (0.009) |
| *Period dummies* | YES | YES | YES | YES |
| Number of observations | 3232 | 3232 | 6300 | 6300 |
| Number of clusters | 120 | 120 | 120 | 120 |

*Notes*: The regressions are based on linear probability models. The binary dependent variable in Model 1 and 2 is the cancellation of the donation in Part 3 in each of the 25 rounds (i.e., the choice of Option 2). These models include only data where the state is bad. The binary dependent variable in Model 3 and 4 takes value one when the decision maker is observing no news. Robust standard errors clustered at group level are in parentheses (each participant count as a distinct group in the INDIVIDUAL treatment). DM for decision maker; *d* for dummy variables. Control variables are: age of the participant (demeaned); gender dummy d(Male); high school grade at the Baccalaureat (BAC) (demeaned); number of past participations in experiments. Period dummies are included with period 1 as the reference category. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; ° $\leq 0.1$.

observed by decision makers about 30% of the time. In the INDIVIDUAL condition, decision makers consume about 6 percentage points more news than in the RANDOM or CHOICE treatments. However, when we regress the consumption of ignorance on treatment dummies and controls for the subjects' preference type, this difference is not statistically significant (see Table 2, column (3)). Again, these results hide heterogeneity between different types of decision makers, which we discuss in Section 4.5.

Does it matter what information decision makers consume? The right panel of 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 all treatments, decision makers systematically choose Option 2 after no news or good news. Since there is no explicit ethical trade-off in these cases, this shows that subjects understand the choices in front of them. By contrast, when decision makers get bad news, only about 40% of their choices are selfish. These fractions are very similar across the three treatments, indicating that conditional on information consumption, types display stable behavior across treatments.

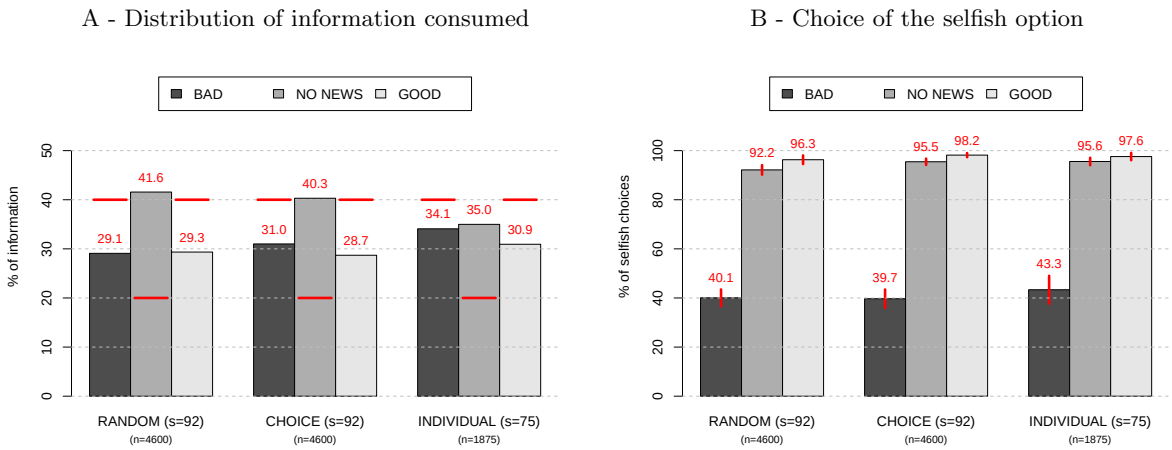


*Figure 2: Information consumption and choices of the decision makers*

*Notes*: Panel A 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 senders. Panel B 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 the group level. Labels below the bars indicate both the number of subjects ($s$) and the total number of choices ($n$).

**Result 1.** *There are no differences across the RANDOM and CHOICE treatments in the aggregate consumption of information and ethical decisions. Subjects in the INDIVIDUAL treatment are more likely to act altruistically.*

## 4.2 Supply of Information

On the supply side, we focus on the suppression of bad news as this is the only news with ethical relevance. As shown in Figure 2, good news does not change the behavior of decision makers.[11] To measure suppression of bad news, we look at the fraction of bad states observed by the sender that were not transmitted to the decision maker in the 25 periods of Part 3. On this measure, senders suppress 27.5% of bad news across both RANDOM and CHOICE. On an individual level, we find that only 1% of senders suppresses all bad news, while 29% of senders suppress no bad news at all. Thus, the large majority of senders suppresses bad news at least sometimes, while about 25% of senders suppress more than half of the bad news they receive. If we apply this metric to individual senders, we can compare the distributions of

[11] Although the suppression of good news does not affect the decision makers' choice, senders may have various reasons to suppress good news. First, given that good news is unlikely to influence the decision or the payoff of the charity, senders may simply like to share cute animal pictures. In Section 5, we show that eliminating cute animal pictures reduces sharing of good news by about 6.5 percentage points, or 20 percent, which is statistically significant. However, even without cute animal pictures, almost 20 percent of good news is suppressed. Figure C.3 shows that symmetric suppression of news is quite common, and we provide more detail in Appendix C. One possible reason for this is to avoid that decision makers infer that "no news means bad news". This suggests that senders think about the inferences that senders are going to make.

sender suppression. Figure C.1 in Appendix C gives an overview of the cumulative distribution of supply choices across treatments. We find no statistical differences between the RANDOM and CHOICE treatments (Kolmogorov-Smirnov test, $p = 0.248$). Appendix C also shows individual examples to illustrate various patterns of information suppression by senders with different suppression strategies.

As discussed in our conjectures, various motives may drive the suppression of bad news. To understand these motives, we first consider the role of beliefs about the demand for ignorance. In Part 1 of the experiment, subjects reported their belief about the number of other participants in the session that preferred not to disclose information. If senders aim to attract more decision makers, these beliefs should inform their disclosure strategies in the CHOICE treatment. Furthermore, if senders want to do decision makers a "favor", then we should also see a correlation between beliefs and suppression in the RANDOM treatment.

Table 3 shows regression evidence to test these motives. In column 1, we regress suppression of information on sender's beliefs, and find a highly significant correlation. In column 2 we introduce a treatment dummy, and in column 3 an interaction of beliefs with the treatment. The interaction term is negative and insignificant, showing that if anything, beliefs play a smaller role in the CHOICE treatment.[12] Thus, rather than strategically increase their chance of being selected by decision makers, these results are consistent with the idea that senders are trying to do decision makers a "favor" by shielding them from difficult trade-offs. In line with this interpretation, beliefs about the demand for information do not predict the suppression of good news, where such trade-offs are absent (see Table G.2). Alternatively, beliefs may also reflect a "false consensus effect" and be a proxy for senders' own preferences for information. If so, our findings could indicate "paternalism": a wish to impose the sender's preferred information or decision on the decision maker.[13]

To see if this projection of preferences explains the effect of beliefs, we control for the sender's preferences in column 4 of Table 3. Although the confidence level and the size of the coefficient decline somewhat, the coefficient on the belief variable remains significant, indicating at most

[12]There are also a number of possible explanations for the lack of a belief effect in the CHOICE treatment. First, it could be due to a different distribution of beliefs across treatments. However, a Kolgomorov-Smirnov test cannot reject the hypothesis of equality of the distributions of beliefs ($p = 0.600$), so this seems unlikely. Second, decision makers could have updated their beliefs during the group phase, on the basis of their experiences. However, repeating the regression analysis presented in Table 3 using only data of the first 5 or 10 periods, yields very similar results. A third explanation is that senders lack the information to be strategic. To reduce complexity, we did not tell senders about the choices of the decision makers nor about the behavior of the other senders. This made it more difficult for senders to optimize their strategy. To speed up learning about demand, senders may have experimented with different strategies, reducing the correlation between beliefs and suppression. We cannot test this explanation within our data-set, but it could be addressed in future research.

[13]In particular, recent results show that many people engage in *ideals-projective paternalism* (Ambuehl et al., 2019), i.e., they restrict others' choices according to their own preferences. Bartling et al. (2020) show that Americans are willing to intervene in the choices of others when it comes to providing information, but less so when it comes to their choices. Our setting providing an intermediate case, as here information is necessary to make an informed decision.

*Table 3: Suppression of bad news by senders*

| | Model 1<br>Est. (S.E.) | Model 2<br>Est. (S.E.) | Model 3<br>Est. (S.E.) | Model 4<br>Est. (S.E.) |
|---|---|---|---|---|
| (Intercept) | 0.114 (0.059)° | 0.114 (0.065)° | 0.078 (0.071) | 0.221 (0.102)* |
| Belief # ignorant | 0.024 (0.006)*** | 0.024 (0.006)*** | 0.031 (0.009)*** | 0.024 (0.009)* |
| d(CHOICE) | — — | -0.000 (0.052) | 0.070 (0.070) | 0.051 (0.069) |
| d(CHOICE)×Belief # ignorant | — — | — — | -0.015 (0.010) | -0.019 (0.010)° |
| Altruistic | — — | — — | — — | -0.139 (0.057)* |
| Avoider | — — | — — | — — | 0.157 (0.064)* |
| $Age - \overline{Age}$ | -0.003 (0.003) | -0.003 (0.003) | -0.003 (0.003) | -0.001 (0.003) |
| *d(Male)* | 0.003 (0.046) | 0.003 (0.045) | -0.004 (0.043) | -0.043 (0.042) |
| $BAC - \overline{BAC}$ | -0.009 (0.011) | -0.009 (0.011) | -0.010 (0.011) | -0.013 (0.010) |
| *# of past participations in exp.* | -0.011 (0.017) | -0.011 (0.017) | -0.009 (0.017) | -0.006 (0.017) |
| *Period dummies* | YES | YES | YES | YES |
| Number of observations | 5389 | 5389 | 5389 | 5389 |
| Number of clusters | 46 | 46 | 46 | 46 |

*Notes*: These regressions are based on linear probability models. The binary dependent variable is the sender's choice to suppress bad news in Part 3 in each of the 25 rounds. 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. Control variables are: age of the participant (demeaned); gender dummy d(Male); high school grade at the Baccalaureat (BAC) (demeaned); number of past participations in experiments. Period dummies are included with period 1 as the reference category. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; ° $\leq 0.1$.

a modest role of projection. This gives some support to the "doing a favor" type of explanation. In addition, both the coefficients for Altruistic and Avoider preferences are significant and sizable, reducing suppression by 13.9 and increasing suppression by 15.7 percentage points, respectively. This last result indicates a role for paternalism, which is not moderated by beliefs about the decision maker's preferences.

Another potential motive for suppression is the presence of cute animal pictures: If decision makers value such pictures, and senders anticipate this, they may try to entertain the decision maker. Part of this should be captured by our belief variable, which measures beliefs about general demand for ignorance. To dig into this further, Section 5 reports the results of treatments that eliminated cute animal pictures. We find very similar results in these treatments. Sharing cute animal pictures reduces the sharing of bad news by about 4 percentage points, but this effect is not statistically significant (see Table G.2, column 1). Cute animal pictures do not seem to be the main driver of suppression of bad news.

To obtain further evidence for senders' motives, we looked at the closing questionnaire, where senders answered the question "According to which principle(s) did you decide to report or not the consequences to the receivers?" With the help of three independent raters, we coded these answers into separate categories. We provide details of the coding scheme and the outcomes in Appendix F. The answers reveal a number of motives, including those discussed

above.[14] Because we have quite a few different motives that have no natural order, the numbers in each category are too small to do reliable statistical analysis on treatment differences. However, in line with common sense, we do see that strategic motives and the wish to help the decision maker are more frequent in the CHOICE treatment, while paternalistic motives and the autonomy of the decision maker are cited more often in the RANDOM treatment.

What was the optimal sender strategy? In Section 4.3, we show that senders who transmit more information than the other two senders in their group are most likely to be chosen by decision makers, followed by senders who are transmit least information. This suggests that the optimal sender strategy is to transmit all information. However, this conclusion is highly dependent on the actual strategies of the other senders in the market as well as on the types of the receivers. As we explained in our design section, we limited information to senders on these aspects to keep the experiment and the analysis tractable, but future research could delve deeper into the strategic motives and interactions of senders.

**Result 2.** *The majority of senders sometimes suppress bad news. Contrary to Conjecture 1, there is no evidence that competitive motives increase suppression in the CHOICE treatment. Suppression is driven by beliefs about the decision makers' preferences and the preferences of the sender.*

### 4.3 Demand for information

We now turn to the demand for information in the CHOICE treatment, where decision makers could choose a sender. On aggregate, we find that decision makers consume information in about the same proportion as it is supplied. In particular, there is no aggregate tendency to seek out or avoid bad news. When we compare the number of bad states seen by the decision makers with the average number of bad states reported to them by the three senders, we find these are almost equivalent: the average ratio between the former and the latter is 1.03 in the CHOICE treatment and 0.99 in the RANDOM treatment.

However, these aggregate result may hide a lot of heterogeneity in information consumption and search strategies. To test Conjecture 2, we examine the likelihood to choose different senders depending on their profile of past information supply, which is available to the decision

[14] As examples of doing favors, some senders justify suppression by mentioning that they aim to "relieve the conscience" of decision makers or "make their decisions easy", and that they tried to "anticipate their expectations". As examples of paternalistic behavior, some suppressed information depending on whether it was "more profitable for them [the decision makers] to know it or not." Senders who sent information frequently cite the importance of giving decision makers a choice to donate and exercise their autonomy, while a few mention the wish to attract more clients.

maker. To summarize the sender's information profile, we rank senders according to the relative level of ignorance they provided to the decision maker in the previous 10 periods, i.e., the ranking in period $t$ is based on the number of times senders disclosed information in periods $t-1$ to $t-10$. Then, we ask how frequently the decision makers chose the sender providing the highest, the intermediate, and the lowest level of ignorance. Note that this approach excludes the first 10 periods from the analysis, as senders have not yet established a history.

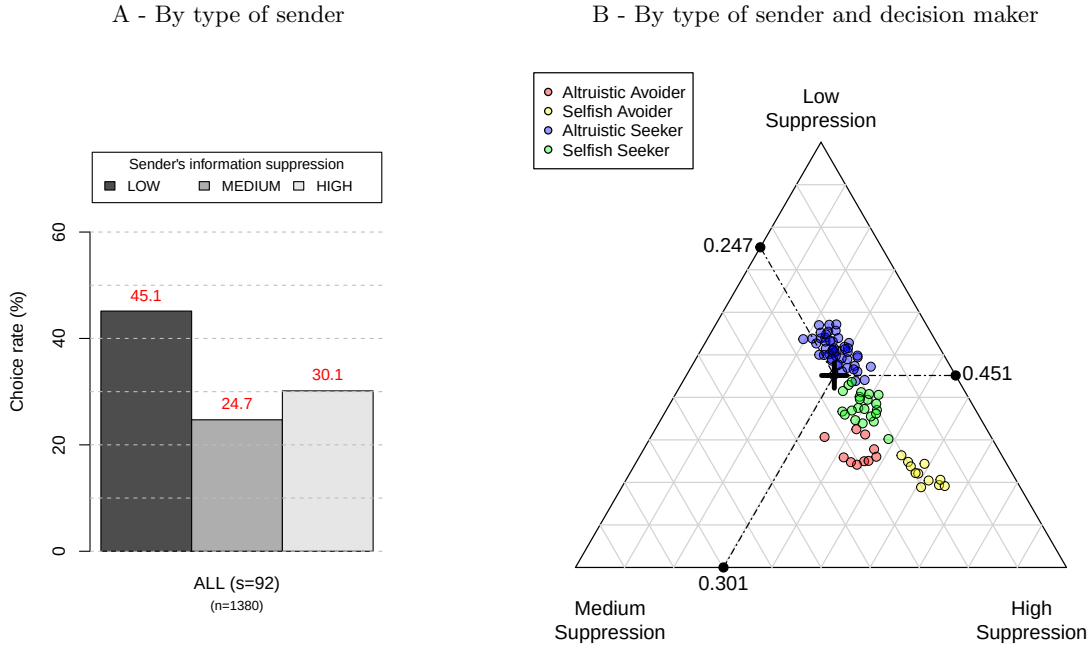


*Figure 3: Predicted probability to choose a given sender*

*Notes*: The figure displays the frequency of choices of the three types of senders in the CHOICE treatment. Senders are ranked (low, medium, and high) according to the relative level of ignorance they provided in the previous 10 periods. Panel A shows displays the predicted probability to choose the sender providing the highest, intermediate, and lowest level of suppression for each of the 92 decision makers. Predicted probabilities are based on Model 1 of Table D.1 of appendix D. Panel B adds the type of the decision maker. The color of the dot captures the type of the decision maker based on decisions in Part 1. The cross shows the average frequency of choice.

Figure 3 shows the frequencies with which different types of senders are chosen. Panel A shows that, on aggregate, the modal choice is the sender that provides the highest level of information. Interestingly, the least popular senders are those with an intermediate profile of information suppression, indicating that some senders are looking for ignorance.

Panel B disagggregates these results by the type of decision maker. It shows a simplex with the predicted probability to choose each sender for each decision maker based on a multinomial logit model—Model 1 in Table D.1 of Appendix D—where the three alternatives are the senders providing relatively Low, Medium, and High ignorance and 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. Predicted probabilities reveal that the aggregate results hide a lot of heterogeneity. The largest group of decision makers (Altruistic Seekers) clearly penalizes non-informative senders: the most informative sender in the group is chosen, on average, 54.6% of the times whereas the least informative sender is chosen, on average, 22.6% of the times. By contrast, a smaller group (Selfish Avoiders) does the opposite: they choose the most informative sender 27.3% of the times and the least informative sender 44.8% of the times. The simplex shows that the effect of heterogeneity is mostly captured by shifting the probability mass from the Low ignorance to the High ignorance sender, while the predicted probability to choose the Medium ignorance sender is about 20-25% and does not change much across types of decision maker. Statistical support for these results is reported in Table D.1 of Appendix D.

Finally, in Appendix D, we provide an additional set of measures of information demand. Most importantly, we quantify the degree of information seeking or avoiding of each individual decision maker by comparing their consumption of bad news to the average amount of such news that is available from senders. As we show in the Appendix, roughly 40% of decision makers consume less bad news than the available average in the market, which indicates information avoidance. The appendix also provides illustrations of different patterns of selection of senders by individual decision makers.

Overall, this analysis supports Conjecture 2 and is summarized in the following result:

**Result 3.** *A majority of decision makers searches for informative senders. However, about one third of the decision makers who avoid information in Part 1 seek out non-informative senders in Part 3.*

**Motives behind avoidance.** While the results in this section show that some subjects deliberately avoid information in the main part of the experiment, this does not pin down their exact motives. Subjects may try to avoid a difficult dilemma, reduce their guilt, or enjoy watching a picture of a cute animal. Pinning down the precise psychological motives is difficult even in individual choice experiments, and has led to an active discussion, see, e.g., Vu et al. (2022). Nevertheless, we can make a number of further comments that add to this discussion.

First, Avoider types who obtain bad news from senders are ready to make a sacrifice for the charity about 53 percent of the times in the RANDOM treatment and about 41 percent of the times in the CHOICE treatment. This suggests that ignorance is not simply borne out of indifference towards the charity. Second, in Section 5 we show that levels of information avoidance are similar in a setting where there are no pictures of cute animals, indicating that

information avoidance is not primarily driven by the entertainment provided by the cute animal pictures. Decision makers in these treatments even mention information avoidance slightly more often than in the treatments with animal pictures. Finally, we also asked subjects in the CHOICE treatments about their motives for selecting a sender. Three independent raters classified the textual responses into several categories, where a text was classified when at least two raters agreed on the category. Table F.2 of Appendix F shows the classification of these responses by treatment. The most commonly reported motive is to select the sender who sends most information (48 percent mention this). The desire to avoid information or see animal pictures is instead mentioned explicitly by about 9 percent of subjects, some of whom mention the desire to make their decision easier or reduce their consciousness. Other motives include randomizing to make sure all senders had a chance to earn something (7 percent), or to reward senders who send information (10 percent). About 8 percent says they used no particular rule, and 17 percent mentioned various other considerations.

### 4.4 Assortative matching

Here, we investigate whether the information demand and supply of senders and decision makers lead to assortative matching between types with similar information preferences, which might help explain the formation of echo chambers. To evaluate assortative matching, we use the two dimensional type classification explained above, based on Part 1 choices. In particular, we look whether decision makers who avoid or seek information in Part 1 are more likely to match with a sender with similar preferences. Figure 4 shows the frequency of choice of senders who avoided disclosing information when deciding in Part 1, by decision maker's type. Frequencies are calculated using data of groups where decision makers have the opportunity to choose either type of sender, i.e., groups where there was at least one Avoider and one Seeker among the senders.

The figure shows that (i) being an Avoider substantially increases the probability of matching with an Avoider-sender and (ii) this effect is stronger for selfish decision makers. Indeed, the probability to choose an Avoider-sender increases by 14 percentage points (from 34.6% to 48.6%) for Altruistic decision makers and by 23.3 percentage points (from 26.0% to 49.3%) for Selfish decision makers. This pattern, which receives statistical support from the regression results reported in models (1) and (2) of Table E.1 in Appendix E, is coherent with the fact that being an Avoider predicts both the demand for ignorance and its supply. The analysis leads to our third result that supports Conjecture 3.

**Result 4.** *In the CHOICE treatment, Avoider-type decision makers match significantly more often with Avoider type senders.*

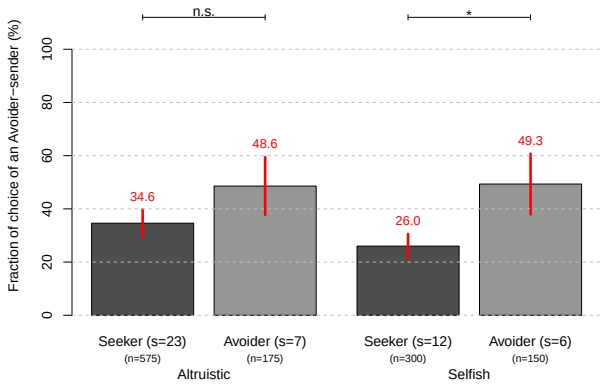


*Figure 4: Assortative matching - choice of an avoider sender by decision maker type.*

*Notes*: The figure displays the frequency of choice of an Avoider sender. Decision makers are split by type, as defined based on decisions in Part 1. The frequencies are calculated on the subset of groups where there is at least one and at most two Avoiders among the senders. Vertical lines represent standard errors based on a linear probability model with clustering at the group level. The two factors on the x-axis and their interaction are the only explanatory variables. Pairwise comparisons reported above the bars are based on a Wald test performed using this estimated model. Signific. codes: * $p \leq 0.05$; n.s. means $p > 0.05$.

## 4.5 Decision maker heterogeneity

How do the treatments affect the information consumption and ethical behavior of different decision makers? In Section 4.1, we already noted that there is no aggregate change in these variables between the CHOICE and RANDOM treatments, whereas subjects behave more altruistically in the INDIVIDUAL treatment. However, these results may hide some interesting variations, as expressed in Conjecture 4. A direct choice of information, and, to a lesser degree, assortative matching in the CHOICE treatment, will lead different types of decision makers to obtain different kinds of information. This, in turn, is likely to lead to diverging decisions between these types. In particular, selfish participants may find it easier to be selfish if they can match with a sender that keeps them uninformed, while altruistic participants can only be altruistic if they find a sender that gives them information. Thus, we expect types to express their preferences more strongly in the CHOICE treatment compared to RANDOM, and, for obvious reasons, even more so in the INDIVIDUAL treatment.

This conjecture is indeed borne out in the data. Panel A in Figure 5 shows the fraction of decisions that cancelled the donation, where we condition on being in the bad state. We split behavior by treatment and type of the decision maker (Selfish vs. Altruistic). As conjectured, Selfish decision makers become about 3 percentage points more selfish in the CHOICE treatment, whereas Altruistic types become 5 percentage points *less* selfish. In the INDIVIDUAL treatment, the divergence between types becomes even more pronounced, as we see a further

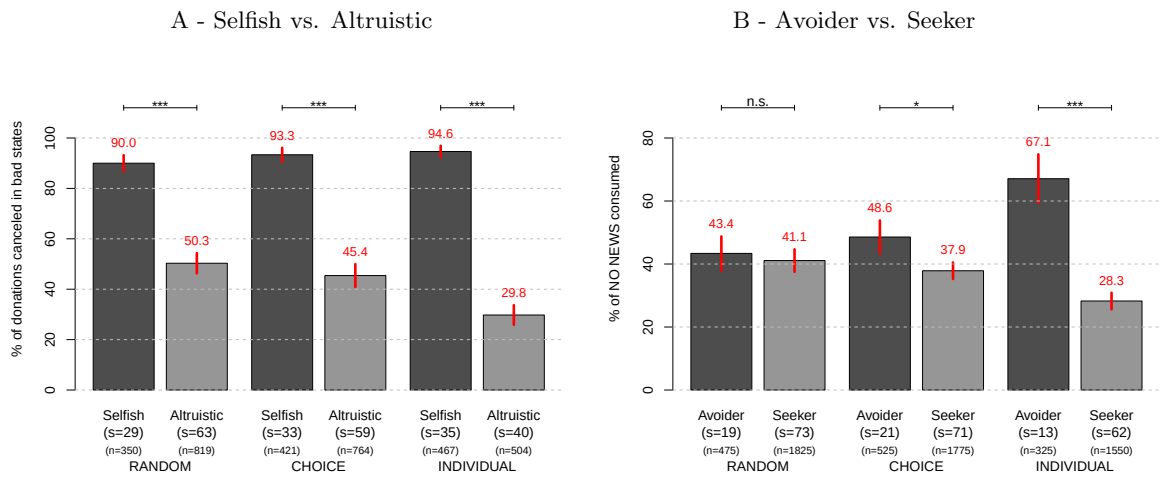


*Figure 5: Selfish decisions in the bad state and ignorance consumption by condition and decision-maker type.*

*Notes*: Panel A displays the fraction of selfish decisions in the different experimental conditions. The fraction of selfish decisions in the RANDOM, CHOICE, and INDIVIDUAL treatments are computed using only the cases where the state is bad. Panel B displays the fraction of NO NEWS consumption in the different experimental conditions. Panel A splits decision makers into Selfish and Altruistic. Panel B splits them into Avoider and Seeker, based on their decision in Part 1. In all panels, vertical lines represent standard errors based on a linear probability model with clustering at the group level. In all models, the two factors on the x-axis and their interaction are the only explanatory variables. Pairwise comparisons reported above the bars are based on a Wald test performed using these estimated models. Signific. codes: * $p \leq 0.05$; ** $p \leq 0.01$; *** $p \leq 0.001$; n.s. means $p > 0.10$.

drop of over 15 percentage points in selfish behavior among Altruistic types. Column 2 of Table 2 confirms this effect statistically, as the dummy for Altruistic types has a smaller coefficient in INDIVIDUAL than in CHOICE and RANDOM. However, the difference-in-difference between the latter two conditions is not statistically significant (Wald test: $\beta = 0.206 - 0.246 = -0.040$, $p = 0.573$).

The result in Panel A also sheds some new light on the increase in aggregate selfishness in RANDOM and CHOICE compared to INDIVIDUAL, that we noticed in Section 4.1. Panel A shows this is driven mostly by the Altruistic subjects. Indeed, Altruistic types are less likely to consume inconvenient information when they depend on an external sender. This may be due to market frictions – Altruists are not always able to match with senders that transfer information – or because Altruists consult uninformative senders as an excuse when given the choice to do so.

In Panel B, we look at the consumption of ignorance, split by Seeker and Avoider types. We see a similar pattern as in Panel A, in that Avoider types become more likely to consume no news in the CHOICE treatment compared to the RANDOM treatment, as they are now less likely to obtain information due to assortative matching with a suppressing sender. The opposite holds for the Seeker type, which is more likely to consume news. Again, the com-

parison with the INDIVIDUAL treatment (left bars) shows that decision-makers' preferences for information matter even more in this case. These results are mirrored in the regressions in Table 2, columns 4 that explain the consumption of no-news. Compared to the INDIVIDUAL condition, the coefficient of Avoider declines in CHOICE and RANDOM. However, the difference-in-difference between these last two conditions is not statistically significant (Wald test: $\beta = -0.231 - (-0.329) = 0.098$, $p = 0.159$). This analysis leads to our fifth result.

**Result 5.** *In line with Conjecture 4, having control over one's information source increases the gap in behavior between different types of decision makers.*

## 5 The role of cute animal pictures

In our main experiments, the alternative to recipient information is to see a cute animal picture (CAP). CAPs mimic the presence of irrelevant distractions of the type that is typical on social media, and also counterbalance the visual information about the recipient that is provided with the payoff information. However, the presence of CAPs may drive the demand for ignorance if participants really enjoy seeing cute animals. It may also drive the suppression of information, if senders anticipate this desire on the side of decision makers, or simply like to share CAPs. To test the impact of CAPs, we ran a "NoCAP" variation of both the RANDOM and CHOICE treatments where we removed the CAPs, while keeping all other aspects the same.

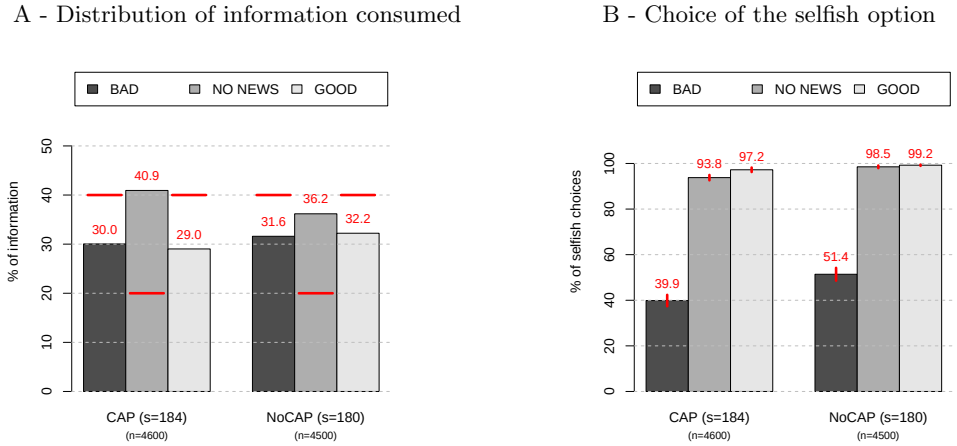


*Figure 6: Information consumption and choices of the decision makers in the CAP and NoCAP treatments*

*Notes*: Panel A displays the distribution of information observed by the decision makers in Part 3, split by treatment. The horizontal lines show the distribution of information the decision makers would observe in case of full disclosure. Panel B 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 ($n$).

## 5.1 Aggregate results

Appendix G provides full results from the NoCAP treatments, following the analysis of our original treatments in the previous section. Here we highlight some key findings, starting with a comparison of the aggregate market results. Since the RANDOM and CHOICE treatments have rather similar results, both in our CAP and in our NoCAP treatments (see Table G.1), we aggregate these two treatments for this analysis. In the NoCAP treatments, 68.7% of decisions in the bad state are selfish and cancel the donation. This is higher than in the original CAP treatments, where the average number was 62.3%. However, as Table G.1 shows, this difference is not significant (column 1), and declines further if we control for the interaction with CHOICE and RANDOM (column 2).

Following our analysis in Section 4.1, Figure 6 disaggregates this result by looking at information sharing and behavior conditional on information. Panel A graphs the average news consumption in the CAP and NoCAP treatment, and shows that the consumption of ignorance is about 5 percentage points lower in the NoCAP treatment. This is in line with the idea that subjects do like to see CAPs, but Table G.1 (columns 3 and 4) shows that the difference is not significant. Panel B shows ethical behavior, conditional on the type of news received. Here we see a rise in selfish behavior among those who consume bad news in the NoCAP compared to the CAP experiment. This is likely a selection effect. The absence of CAPs makes recipient information relatively more attractive, so the sample who consumes bad news is larger in the NoCAP treatment. However, these marginal consumers are less altruistically motivated and hence, more likely to act selfishly.

## 5.2 Supply and demand

Appendix G shows the NoCAP equivalents of all figures and tables presented in the above analyis of supply and demand. All main patterns go through qualitatively, with only minor differences between the CAP and NoCAP treatments. Nevertheless, a few patterns are worth pointing out. First, Avoider preferences are stronger predictors of both senders' choices to send information and of receivers' choices for senders in the NoCAP than in the CAP experiment. This is likely due to a similar selection effect as we mentioned above. As CAPs were absent in Part 1 where we classify the preferences of subjects, those classified as Avoiders in the NoCAP treatment may have stronger preferences for ignorance than Avoiders in the CAP treatment. In fact, the sample of Avoiders in the NoCAP treatment is relatively small, and this prevents us from running the assortative matching analysis in Section 4.4 for these treatments.

Second, while the overall suppression of bad news in NoCAP is a few percentage points lower than in CAP (23.8% vs. 27.5%), we still see substantial suppression. Comparing our regression Table 3 with Table G.4, we see similar patterns. In particular, sender's beliefs about the information preferences of decision makers remain highly predictive of their choice to send

information, both in the RANDOM and CHOICE treatments. Sender motives are similar across CAP and NoCAP treatments. To compare the effect of CAPs on the suppression of good and bad news, in Table G.2 we separately regress the suppression of bad and good news on a treatment dummy for the NoCAP treatments and sender characteristics. The NoCAP treatments reduce the suppression of both types of news, but the effect is larger for good news (6.5 vs. 4.6 percentage points) and statistically significant only for that case. Thus, there is some evidence that senders aim to please decision makers with CAPs, but this mostly affects suppression of good news with no consequences on charity payoffs.

Third, we find that decision makers select senders in very similar ways in the CAP and NoCAP treatments. Comparing Figure 3 with Appendix Figure G.2 shows that the fractions of decision makers' choice of the least and most informative senders are almost identical across the CAP and NoCAP CHOICE treatments. This indicates that CAPs are not the main reason people avoid information about the beneficiary. This leads to our last result.

**Result 6.** *The slightly higher consumption of informative news in the NoCAP treatments does not induce a more ethical behavior. CAPs do not appear to be the key driver of suppression and avoidance of inconvenient information.*

## 6 Conclusion

We have shown how social interactions can produce willful ignorance through the behavior of both sides of the interaction. First, senders are willing to share irrelevant distractions instead of relevant information, a decision driven both by their own preferences and their beliefs about the sender's preferences. Second, about one third of decision makers prefer to "shoot the messenger" of inconvenient information, and actively seek out uninformative senders. These behaviors result in more (voluntary and involuntary) selfish decisions than when people act in isolation. Thus, we show that the phenomenon of information avoidance that has been documented in individual decisions extends to situations where people depend on others to obtain information.

The experiment also shows how social interactions produce differences in information between different people. When subjects can choose their own senders, we find assortative matching between information-avoiding decision makers and information-suppressing senders. Behavior in this setting is also closer to that observed in an individual condition, where people have full control over their information. These comparisons show the importance of preference-based sorting mechanisms for the information consumption and subsequent behavior of different individuals.

These results have a number of applications outside the lab. First, our setting can be seen as a stylized social media platform, where people follow others based on their information profile. Even though subjects do not communicate directly about their preferences, our results show similar dynamics of assortative matching as have been found on social media, confirming the "filter bubbles" phenomenon (e.g., Aiello et al., 2012). As long as people can choose their own connections, they tend to behave in a homophilous way by selecting like-minded sources of information. This also echoes studies on homophily in the endogenous selection of peers in the moral domain, both with (Gross et al., 2018) and without (Charroin et al., 2021) complementarity between peers. A second application is in organizational design, relating to the relationships between an executive and a consultant. The experiment shows that having decision makers be guided by independent external advisers is not a guarantee for more ethical choices, as advisers may suppress inconvenient information either to please the decision maker or to impose their own agenda in a paternalistic perspective. It also shows that institutional details affect the balance of power. We find that decision makers' preferences are more predictive of their ethical decisions when they can choose their own advisers. Giving the executive the power to choose advisers puts a greater onus on the executive's character, while strengthening the independence of the bureaucracy does the opposite, but without leading to more ethical decisions.

Our experiment provides a starting point for further investigations of the transmission of inconvenient information. For instance, as we discussed above, our design suggests that senders' decisions result from a complex mix of motives. Given how often we make the decision to share (or not) content with others both on and offline, this is an important area for future research. Furthermore, the design of our interactions could be extended in several directions. What happens when senders' information is known to be more or less noisy? Or when senders can actively falsify, rather than just suppress information, engaging in cheap talk or "fake news" generation rather than disclosure? What if decision makers can consult multiple senders for a second opinion? And what if senders have to bear some accountability for the consequences of the decision makers' choices? Even if these questions are too numerous to answer in a single paper, our framework could prove useful to investigate them in the future.

## 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.

Ambuehl, S., B. D. Bernheim, and A. Ockenfels (2019): "Projective paternalism," *NBER Working paper 26119*.

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.

Bartling, B., C. A. W., H. Hermes, M. Skiveness, and B. Tungodden (2020): "Free to Fail: Paternalistic Preferences in the U.S." *Mimeo obtained via personal interaction.*

Bartling, B., R. A. Weber, and L. Yao (2015): "Do markets erode social responsibility?" *The Quarterly Journal of Economics*, 130, 219–266.

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.

Charness, G. and M. Sutter (2012): "Groups make better self-interested decisions," *Journal of Economic Perspectives*, 26, 157–176.

Charpentier, C. J., E. S. Bromberg-Martin, and T. Sharot (2018): "Valuation of knowledge and ignorance in mesolimbic reward circuitry," *Proceedings of the National Academy of Sciences of the United States of America*, 115, E7255–E7264.

Charroin, L., B. Fortin, and M. C. Villeval (2021): "Homophily, Peer Effects and Dishonesty," *GATE Working Paper*, 2021-07.

Chen, Z. C. and T. Gesche (2016): "Persistent bias in advice-giving," *Working Paper.*

Coffman, L. C. and A. Gotthard Real (2019): "Moral Perceptions of Advised Actions," *Management Science*, 65, 3904–3927.

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., D. M. Cain, and R. M. Dawes (2006): "What you dont know wont hurt me: Costly (but quiet) exit in dictator games," *Organizational Behavior and Human Decision Processes*, 100, 193–201.

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.

Espinosa, R. and J. Stoop (2019): "Do People Really Want to Be Informed? Ex-ante Evaluations of Information-Campaign Effectiveness," *Mimeo.*

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.

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

Foerster, M. and J. J. Van der Weele (2021): "Casting doubt: Image concerns and the communication of social impact," *Economics Journal*, 131, 28872919.

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 (2020): "Ambiguity and excuse-driven behavior in charitable giving," *European Economic Review*, 124.

Gneezy, U., S. Saccardo, M. Serra-Garcia, and R. van Veldhuizen (2020): "Bribing the Self," *Games and Economic Behavior*, 120, 311–324.

Gross, J., M. Leib, T. Offerman, and S. Shalvi (2018): "Ethical Free Riding: When Honest People Find Dishonest Partners," *Psychological Science*, 29, 1956–1968.

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.

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

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.

Larson, T. and M. C. Capra (2009): "Exploiting moral wiggle room: Illusory preference for fairness? A comment," *Judgement and Decision Making*, 4, 467–474.

Lazear, E. E. P., U. Malmendier, and R. A. Weber (2012): "Sorting in experiments with application to social preferences," *American Economic Journal: Applied Economics*, 4, 136–163.

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.

Momsen, K. and M. Ohndorf (2019): "Information avoidance, selective exposure, and fake (?) news: A green market experiment," Tech. rep., Working Papers in Economics and Statistics.

——— (2020): "When do people exploit moral wiggle room? An experimental analysis in a market setup," *Ecological Economics*, In Press.

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.

Schram, A. and G. Charness (2015): "Inducing social norms in laboratory allocation choices," *Management Science*, 61, 1531–1546.

Serra-Garcia, M. and N. Szech (2019): "The (In)Elasticity of Moral Ignorance," *Mimeo Karlsruhe Institute for Technology*.

Sharot, T. and C. R. Sunstein (2020): "How people decide what they want to know," *Nature Human Behavior*, 4, 14–19.

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

Vu, L., I. Soraperra, M. Leib, J. J. V. der Weele, and S. Shalvi (2022): "Willful Ignorance: A Meta Analytic Review," Unpublished manuscript, available on request, University of Amsterdam.

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 for CAP treatments were translated from French. Instructions for the No-CAP treatments are into brackets.

## A.1 Instructions for the RANDOM 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. [*Last sentence omitted in the NoCAP treatment*]

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

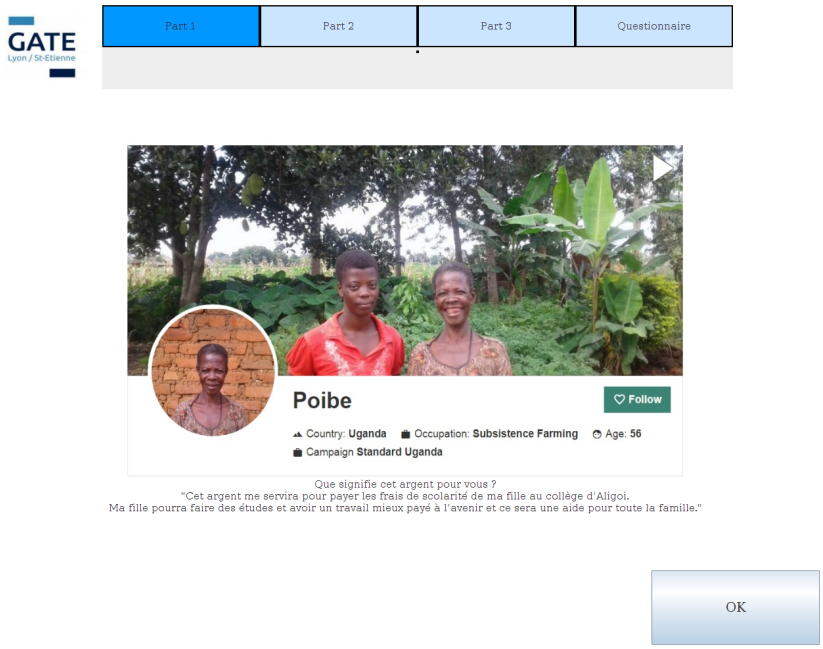


the donation, just like in decision 1. With 50 chances out of 100 the situation is reversed, so choosing "Option 1" cancels the donation and "Option 2" confirms the donation. You are not informed which situation is chosen by the program, and the consequences for GiveDirectly are replaced by "???".

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" [ "No information"]. This means that your screen will display the picture of a cute animal, as illustrated in the screenshot below. [ "This means that you will not see any picture on your screen"] 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" [ "No information"]. 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" [ "No information"] and reported your guess, you will have to choose between "Option 1" and "Option 2". Your earlier choice between "Recipient" and 'Cute animal" [ "No information"] 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" [ "No information"].
4. You report your guess about the numbers of other participants in the session who chose 'Recipient" or "Cute animal" [ "No information"].
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.*

# 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 [ "No information"] 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" [ "No information"] , the Receiver will see the picture of a cute animal but not the consequences for GiveDirectly [ "will not see any information nor 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" [ "No information"] :

- 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" [ "No information"] 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" [ "No information"] .
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.

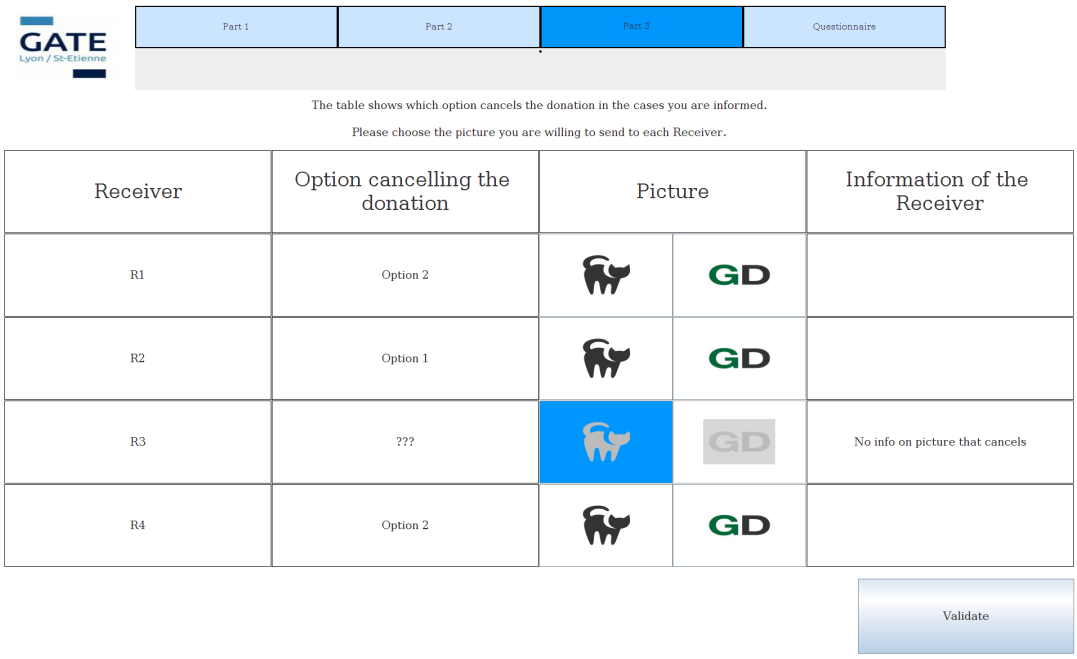


## Choice of the Sender

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 [ "No information", no information 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 [ "No information"] 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 [ the symbol $\emptyset$] 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" [ "No information"]. 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 [ $\emptyset$ for "No information"] or the symbol GD for a recipient of GiveDirectly). The symbol of an animal [ $\emptyset$] 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

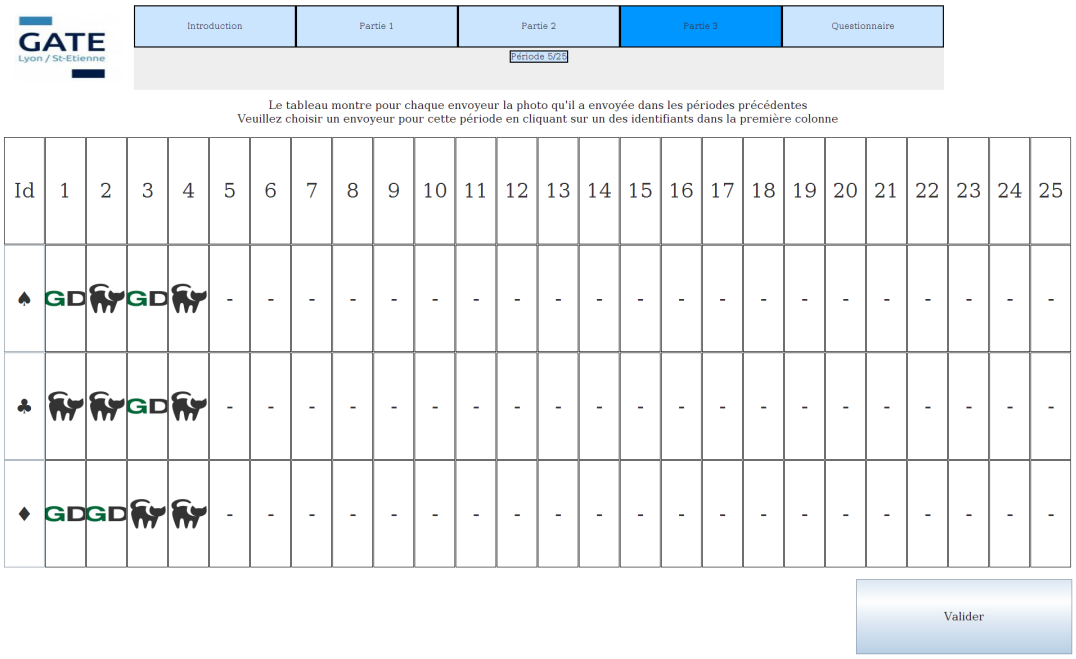


in the group are only visible to the Receivers, not to the other Senders. In this example, symbols have been chosen randomly.

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" [ "No information"] ) 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" [ "No information"] ).
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" [ "No information"] ) 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.[*Last sentence omitted in the NoCAP treatment*]

*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 CHOICE 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.[*Last sentence omitted in the NoCAP treatment*]

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.

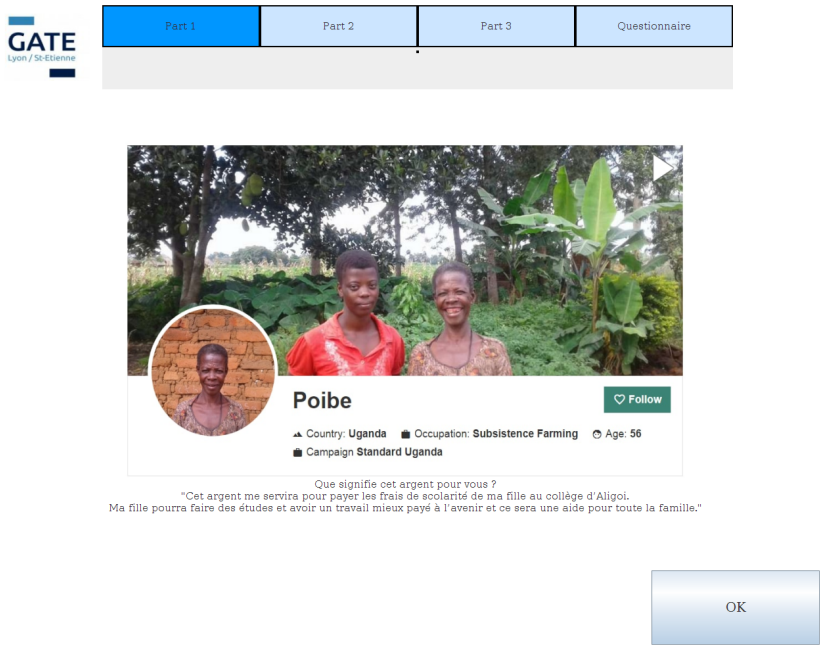


For **Decision one** you will see on your screen before your choice the picture and testimonial of a potential recipient of the donation, who has passed the screening by GiveDirectly, as illustrated in the screenshot below. Choosing "Option 1" will result in €9 for yourself and will *not cancel* the 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" ["No information"]. This means that your screen will display the picture of a cute animal, as illustrated in the screenshot below ["This means that you will not see any picture on your screen"]. 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" ["No information"]. 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"["No information"] and reported your guess, you will have to choose between "Option 1" and "Option 2". Your earlier choice between "Recipient" or 'Cute animal" ["No information"] 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" [ "No information"] .
4. You report your guess about the numbers of other participants in the session who chose 'Recipient" or "Cute animal" [ "No information"] .
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.*

# 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" [ "No information"], the Receiver will see the picture of a cute animal, but not the consequences for GiveDirectly [ "will not see any information nor 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" [ "No information"]:

- 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" [ "No information"] 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" [ "No information"].
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.

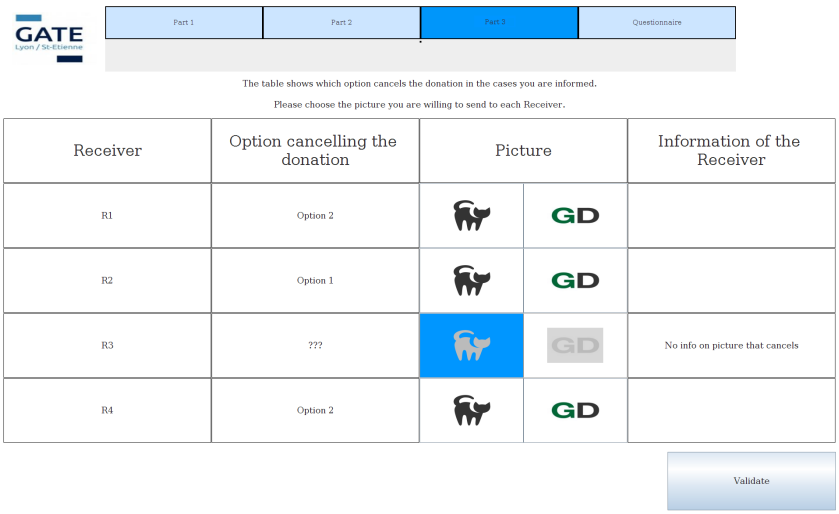


In each period, the program randomly determines the consequences of each option for GiveDirectly, independently for each Receiver. Thus, these consequences can differ across periods and across Receivers.

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" [ "No information"] the picture of a cute animal is [this will be] 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 [ "No information"] 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 [the symbol $\emptyset$] 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" [ "No information"] . 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 [the symbol $\emptyset$ for "No information"] or the symbol GD for a recipient of GiveDirectly). The symbol of an animal [$\emptyset$] reflects *either* that the Sender had no information,

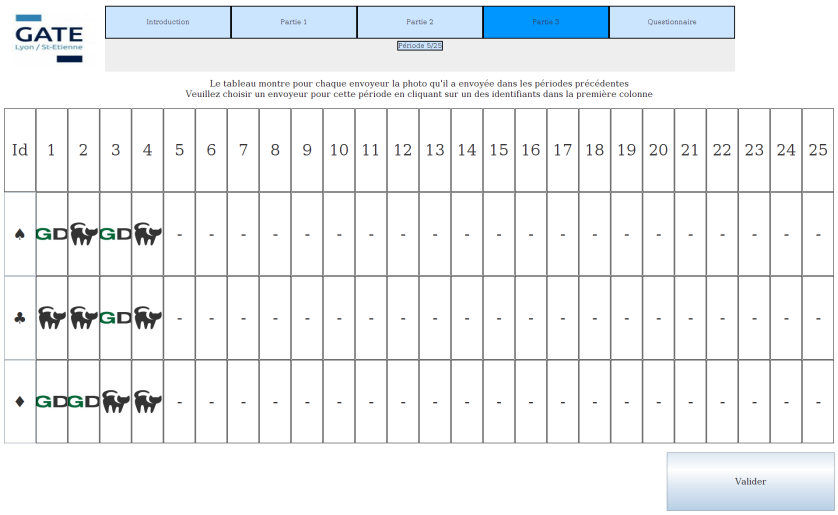


*or* that the Sender received the information and decided to share this picture [not to share it; the Receiver cannot distinguish between these possibilities. The past choices of the Senders in the group are only visible to the Receivers, not to the other Senders. In this example, symbols have been chosen 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" ["No information"]) 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" ["No information"]).
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" ["No information"]) 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.[*Last sentence omitted in the NoCAP treatment*]

*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.3 Instructions for the INDIVIDUAL 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.

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

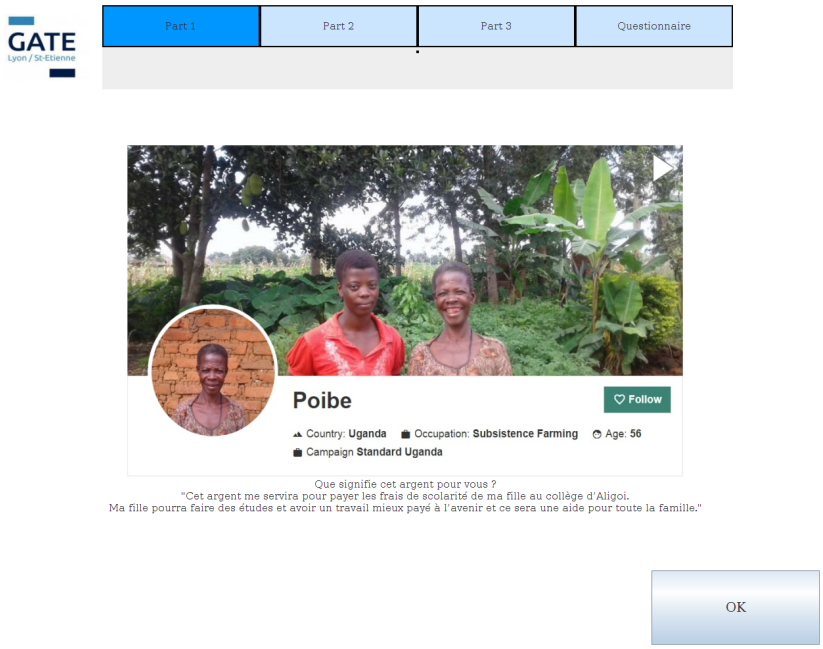


determines randomly which one of the two options will result in a cancellation of the donation of €15 to Give Directly. With 50 chances out of 100, choosing "Option 2" cancels the donation while "Option 1" confirms the donation, just like in decision 1. With 50 chances out of 100 the situation is reversed, 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" and '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".

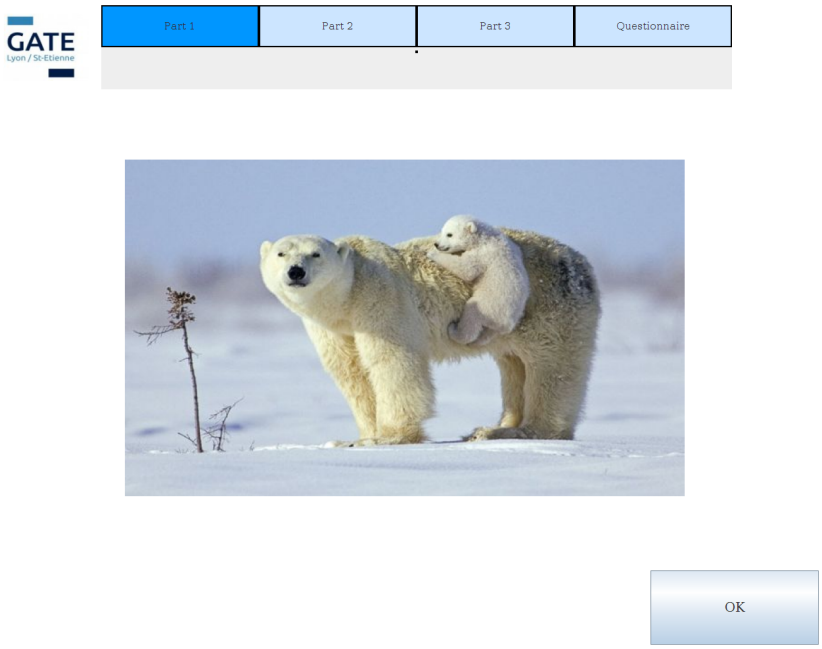


5. Your screen displays the information you chose in step 3, and you choose between Option 1 and Option 2.

*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.*

# Part 2

This part has 25 periods. We now describe each of these periods. In each period, you have to choose between "Option 1" and "Option 2". The consequences from this decision are the same as before:

- "Option 1" results in €9 for you and "Option 2" results in €15 for you.
- 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.

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

You are not informed of the consequences of each option for GiveDirectly. However, before choosing between "Option 1" and "Option 2" you can obtain information from the program.

With 80% chance, the program makes available the information on the consequences of each option for GiveDirectly chosen by the program for the current period. You have to decide which set of information to receive. As before, if you choose "Recipient", i) you will be informed on which option cancels the donation before choosing an option, and ii) the picture of a typical recipient with his/her testimonial will be displayed on your screen. If you choose "cute animal", the picture of a cute animal will be displayed on your screen and you will not learn which option cancels the donation, neither before nor after your choice of option.

With 20% chance the program makes available the information on the consequences of each option and the picture of a cute animal is displayed automatically on your screen.

## Summary of the Decisions

1. The program randomly selects which option cancels the donation.
2. With 80% chance, the program makes available the information on the consequences of each option. If information is available, you have to choose the set of information to receive ("Recipient" or "Cute animal".

3. You decide between Option 1 and Option 2.
4. The same procedure applies for the 25 periods.

## Summary of Earnings

The program randomly selects one of the 27 periods for payment (the two decisions in part 1 and the 25 periods in part 2). For the randomly selected period:

- **You** earn either €9 or €15, depending on the chosen option in that period.
- **GiveDirectly** will receive a donation of €15 if the choice in the selected period does not cancel the donation.

## End of the Session

At the end of part 2 a questionnaire will be displayed on your screen and then you will receive a feedback on your earnings in the session.

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

# B Appendix: Additional Tables and Figures

## Summary of the sessions

*Table B.1: Summary of sessions*

| Session | Treatment | N Participants | Females | Mean Age | Mean Payoff |
|---|---|---|---|---|---|
| 1 | CAP-CHOICE | 21 | 66.67% | 29.10 | 18.76 |
| 2 | CAP-CHOICE | 21 | 61.90% | 23.38 | 18.52 |
| 3 | CAP-CHOICE | 21 | 52.38% | 23.67 | 18.81 |
| 4 | CAP-CHOICE | 21 | 52.38% | 23.43 | 19.33 |
| 5 | CAP-CHOICE | 21 | 52.38% | 21.38 | 18.33 |
| 6 | CAP-CHOICE | 21 | 57.14% | 20.24 | 18.76 |
| 7 | CAP-CHOICE | 21 | 38.10% | 21.52 | 19.05 |
| 8 | CAP-CHOICE | 14 | 42.86% | 20.93 | 18.21 |
| 9 | CAP-RANDOM | 21 | 40.00% | 21.48 | 18.86 |
| 10 | CAP-RANDOM | 21 | 80.95% | 20.76 | 17.67 |
| 11 | CAP-RANDOM | 21 | 57.14% | 20.62 | 18.48 |
| 12 | CAP-RANDOM | 21 | 57.14% | 23.52 | 15.48 |
| 13 | CAP-RANDOM | 21 | 61.90% | 23.43 | 18.24 |
| 14 | CAP-RANDOM | 21 | 61.90% | 21.67 | 18.90 |
| 15 | CAP-RANDOM | 14 | 57.14% | 20.79 | 19.57 |
| 16 | CAP-RANDOM | 21 | 38.10% | 23.10 | 19.14 |
| 17 | NoCAP-RANDOM | 21 | 52.38% | 21.24 | 19.38 |
| 18 | NoCAP-CHOICE | 21 | 66.67% | 20.52 | 19.90 |
| 19 | NoCAP-RANDOM | 21 | 38.10% | 20.76 | 19.86 |
| 20 | NoCAP-CHOICE | 21 | 57.14% | 21.29 | 20.57 |
| 21 | NoCAP-RANDOM | 21 | 47.62% | 23.76 | 20.24 |
| 22 | NoCAP-CHOICE | 21 | 60.00% | 21.14 | 20.19 |
| 23 | NoCAP-RANDOM | 21 | 52.38% | 25.00 | 19.43 |
| 24 | NoCAP-CHOICE | 21 | 61.90% | 22.52 | 19.43 |
| 25 | INDIVIDUAL | 10 | 20.00% | 21.40 | 19.70 |
| 26 | NoCAP-CHOICE | 21 | 71.43% | 20.48 | 19.57 |
| 27 | INDIVIDUAL | 14 | 71.43% | 21.00 | 19.21 |
| 28 | NoCAP-RANDOM | 21 | 57.14% | 21.33 | 19.48 |
| 29 | NoCAP-CHOICE | 21 | 61.90% | 20.52 | 18.81 |
| 30 | NoCAP-RANDOM | 21 | 47.62% | 21.00 | 19.00 |
| 31 | NoCAP-CHOICE | 21 | 57.14% | 20.52 | 19.81 |
| 32 | INDIVIDUAL | 17 | 41.18% | 24.65 | 17.94 |
| 33 | NoCAP-RANDOM | 14 | 35.71% | 23.36 | 19.43 |
| 34 | NoCAP-CHOICE | 14 | 64.29% | 21.00 | 19.86 |
| 35 | NoCAP-RANDOM | 14 | 71.43% | 21.64 | 20.07 |
| 36 | INDIVIDUAL | 20 | 70.00% | 21.65 | 18.30 |
| 37 | INDIVIDUAL | 14 | 57.14% | 21.57 | 18.86 |
| Total | — | 712 | 55.63 | 22.07 | 19.04 |

*Notes*: The table reports the number of participants, the percentage of females, the mean age of the participants, and the mean participant's payoff in Euros, per session. The smaller number of participants in some 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.

# C Appendix: Additional Figures on supply choices

In this Appendix, we show more details of the distribution of supply, using our metric for the suppression of bad news, which we call the $s$ statistic, defined as

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

We also provide additional figures displaying the roles of senders' beliefs and senders' suppression of good news.

## Supply: Overview of suppression of bad news

Figure C.1 shows the cumulative distribution of $s$-statistics over senders 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, almost one third of senders in each treatment transmits all bad news and the large majority suppresses at least some news. Third, about 25% of senders suppress more than half of the bad news they receive ($s > 0.5$). Appendix C gives individual examples to illustrate various patterns of information suppression by senders with different $s$-statistics.

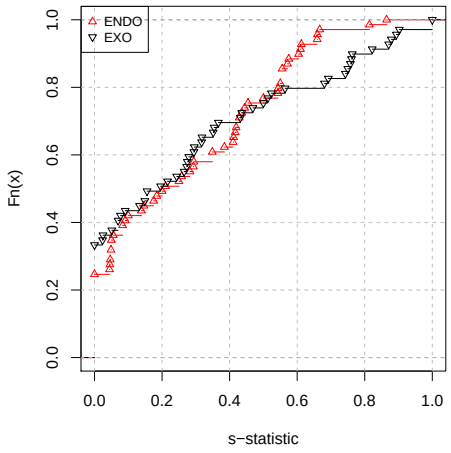


*Figure C.1: Suppression of bad news by senders*

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

## Supply: Correlation between beliefs and suppression of bad news

The left panel of Figure C.2 shows the suppression rates split by senders who believe that less than 50% of other participants in the session prefer animal pictures over recipients' info ("belief low") and senders who believe the fraction is 50% or more ("belief high"). In both treatments, the (minority of) senders who believe the majority prefers distractions are more likely to suppress information. However, contrary to a strategic motive of attracting clients, this difference is larger in the RANDOM treatment than in the CHOICE treatment.

These results are consistent with the idea that senders are trying to do decision makers a "favor", rather than strategically adjust their behavior to increase their chance of being selected by decision makers. However, beliefs may also reflect a "false consensus effect" and be a proxy for senders' own preferences for information. If so, our findings could indicate "paternalism": a wish to impose the sender's preferred information or decision on the decision maker.

To investigate this possibility, we consider the correlation between sender beliefs and sender preferences. The right panel of Figure C.2 shows the beliefs split by the senders' own preferences in the dilemma, based on the individual decisions in Part 1. In line with the "projection" hypothesis, senders who prefer to avoid receiver information ("Avoiders") are more likely to think that others prefer to do so. The impact of taking a selfish decision in the dilemma ("Selfish") is less clear.

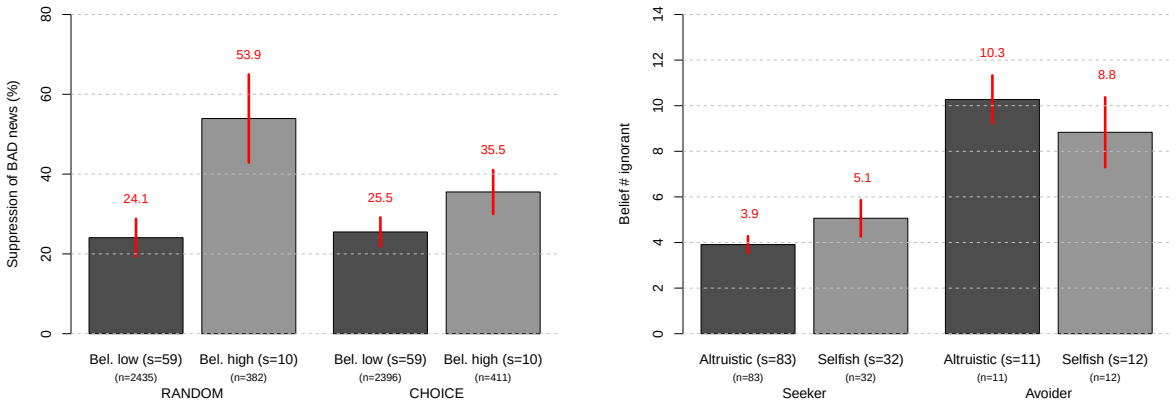


*Figure C.2: Sender beliefs and suppression of bad news*

*Notes*: The left panel shows the impact of sender's beliefs about the number of subjects (in the session) that want to remain uninformed on the rate of suppression of bad news. Belief high (low) is for senders that believe the majority wants to remain ignorant (wants to know). The right panel shows the average belief about the number of subjects that want to remain ignorant conditional on the sender's preferences as revealed in Part 1 of the experiment. Vertical bars show the standard errors. In the left panel, 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 3. Compared to the table, the model only includes the treatment dummy, a dummy that is equal to 1 when the sender believes that the majority wants to remain uninformed, and their interaction as explanatory variables.

## Supply: Suppression of good news

Although the suppression of good news does not affect the decision makers' choice, senders may suppress information symmetrically to avoid that decision makers infer that "no news means bad news". To investigate this, Figure C.3 shows each sender's suppression pattern for both good and bad news. We observe that most of the observations are close to the diagonal, indicating symmetric suppression. Fisher tests reject the null hypothesis that the senders have the same faction of suppression for bad and good news for only 35 out of 138 senders. Of these, 21 systematically suppress bad news more than good news and 14 suppress good news more than bad news. Thus, only a minority systematically suppresses news asymmetrically, which is consistent with the idea that senders think about the (Bayesian) inferences of the decision maker.

At first glance, it is somewhat puzzling that 14 senders suppress good news more than bad news. One rationale for suppressing mostly good news is to signal information suppression in order to lure decision makers who would prefer to remain ignorant. Indeed, such a strategy makes the picture of the beneficiary (and potential victim) more salient when news are bad. For example, one subject reported in the final questionnaire: "When option 1 cancelled the donation, regardless of whether I give information or not, the Receiver would choose to earn 15 Euros (option2), which is good since it gives money to the charity. Even if it has no impact on their decision, I prefer send them a picture of a cute animal. When option 2 cancelled the donation, I sent the picture of GD hoping that this would force the receivers to think further, so that the charity would receive a donation (option 1), even if they would earn only 9 Euros." By providing them with bad news, an altruistic sender may exert pressure on such decision makers to act altruistically, thus increasing revenue for the charity. This idea is consistent with the data: of the 14 decision makers who mostly suppress good news, 12 are Altruistic and 2 are Selfish. Again, this suggests that senders employ rather sophisticated strategies.

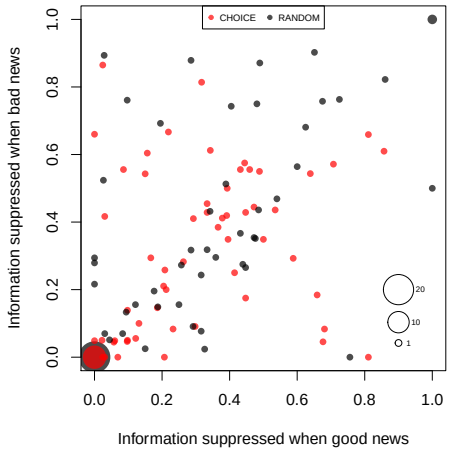


*Figure C.3: Individual suppression rates of good and bad news by treatment*

*Notes*: The figure displays suppression of good news (x-axis) and bad news (y-axis). Each dot is a sender. Multiple senders with the same suppression patterns are indicated by bigger dots.

## Supply: Illustrations of the supply of information by senders with various $s$-statistics

The following three figures correspond to three senders with different $s$-statistics. In the figures, each line corresponds to one of the four decision makers in the sender'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 sender 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.

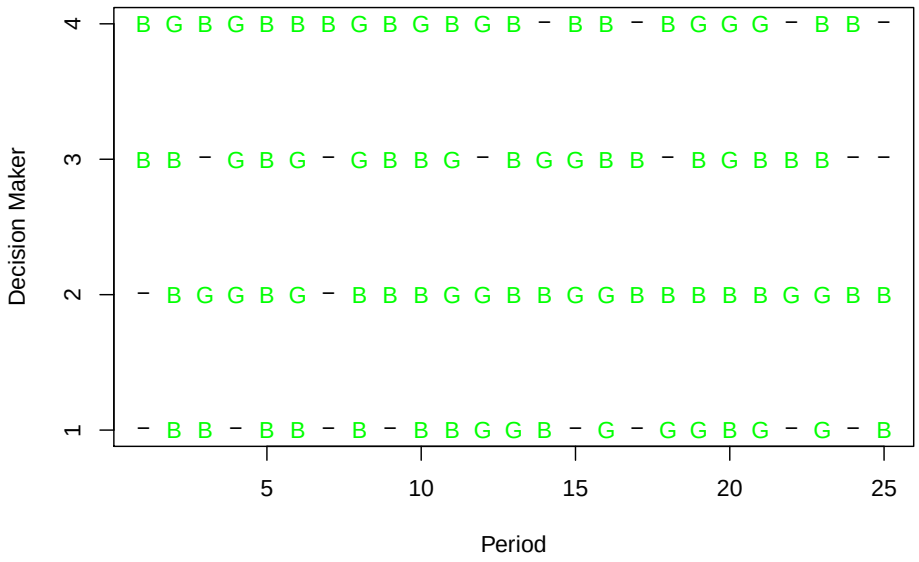


*Figure C.4: Transmission of information to decision makers by a sender with s-statistic = 0*

*Notes*: Figure C.4 illustrates the case of a sender that transmits all information ($s$-statistic = 0).

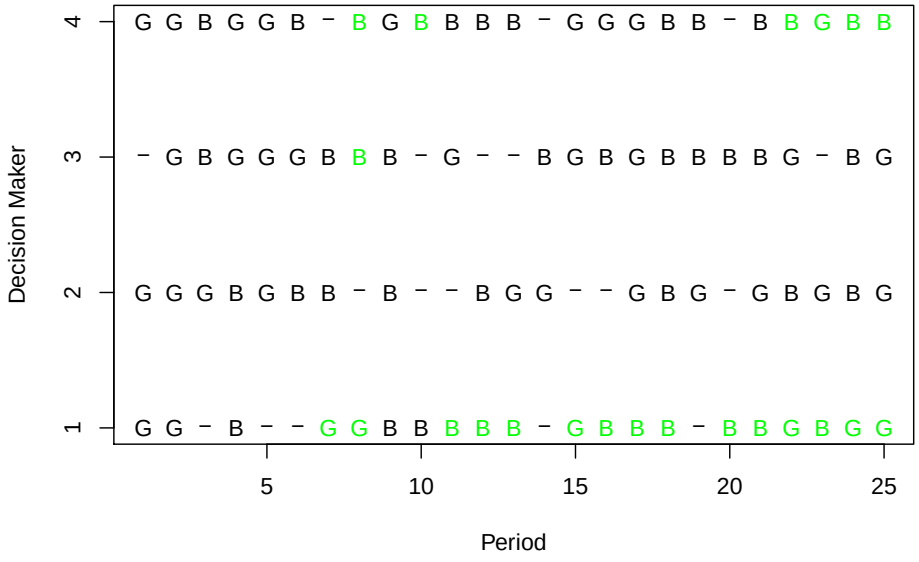


*Figure C.5: Transmission of information to decision makers by a sender with s-statistic = 0.66*

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

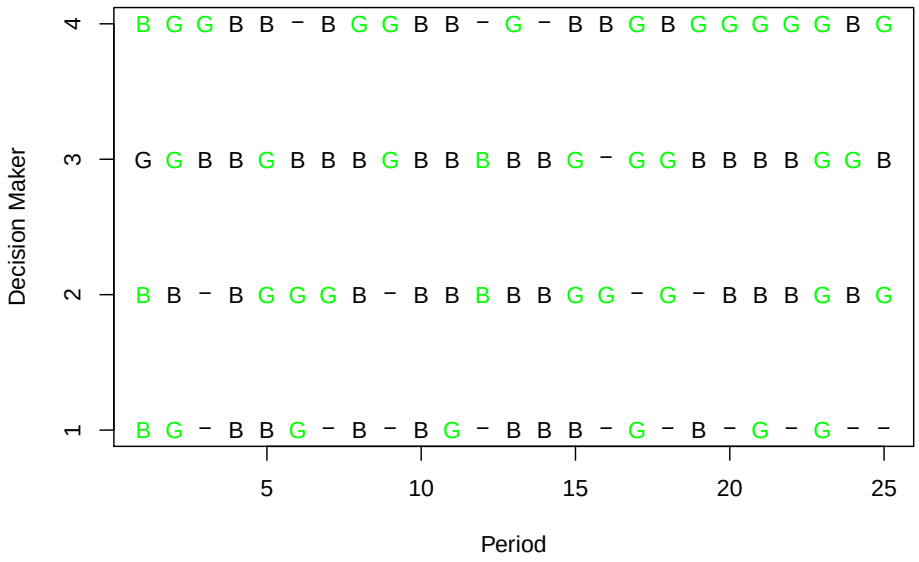


*Figure C.6: Transmission of information to decision makers by a sender with s-statistic = 0.89*

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

# D Appendix: Additional analysis of the demand for information

In this Appendix, we provide additional figures and tables to understand information demand in the main CHOICE treatment (with cute animals pictures), and use a number of additional methods to classify decision makers' selection strategies.

## Probability to choose the senders — Multinomial models

Table D.1 reports multinomial logit models where the three alternatives are the senders 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 D.1 show the effect of the decision maker 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 sender that provides the lowest level of ignorance compared to the odds to choose the sender that provides the intermediate level (Medium ignorance equation), and compared to the odds to choose the sender that provides the highest level of ignorance (High ignorance equation). Results are robust to the inclusion of fixed effects at the group level.

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

| | Model 1<br>Est. (S.E.) | Model 2<br>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)* |
| $\sigma_M$ | 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)* |
| $\sigma_H$ | 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 senders ordered by the amount of ignorance supplied in the previous 10 periods (the baseline alternative is the sender 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. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; ° $\leq 0.1$.

To help interpret these findings, Figure D.1 shows the predicted probability to choose each sender 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 sender, while the predicted probability to choose the Medium ignorance sender 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 senders.

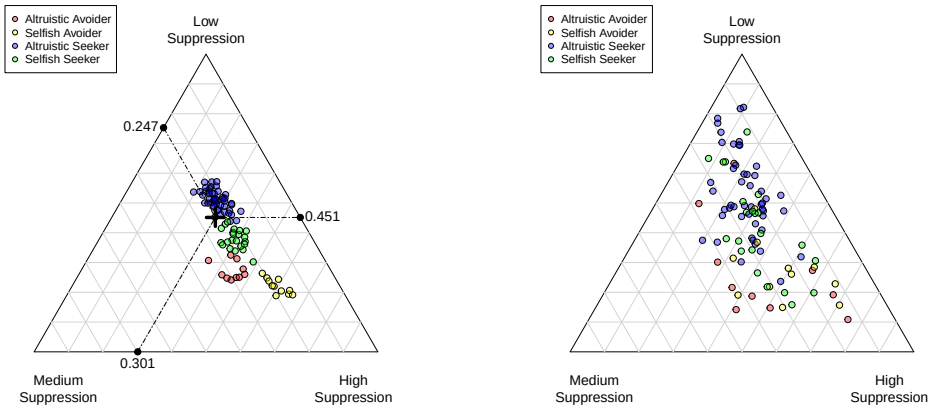


*Figure D.1: Predicted probability of the decision makers' choice of sender*

*Notes*: The figures display the predicted probability to choose the sender 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 D.1 and the figure on the right reports predictions based on Model 2 of Table D.1. The color of the dots captures the type of the decision maker elicited in Part 1. The cross shows the average frequency of choice.

## Demand as measured by news consumption

Here, we compute how often a decision maker actually observes bad news, compared to the bad news that is available from senders in the group. 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 senders he or she was matched with throughout the 25 periods of Part 3:

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

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

Figure D.2 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 (43.5%) decision makers have a $d < 1$, and can be classified as information avoiders in Part 3, 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 is a noisy measure since groups differ in the distribution of news, and hence in the possibility to become more or less informed. In an extreme case where all senders 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.[15] 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

[15]The procedure is as follows: (i) keeping fixed the senders' behavior, we simulate the choice of each decision maker in each period under the assumption that he/she randomly selects one of the 3 senders; (ii) given the simulated choices of the decision maker, we compute the implied $d$-statistic; (iii) we repeat the procedure 100000 times. The confidence intervals are obtained by taking the 5th and 95th percentiles of the simulated $d$-statistic over the 100000 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.

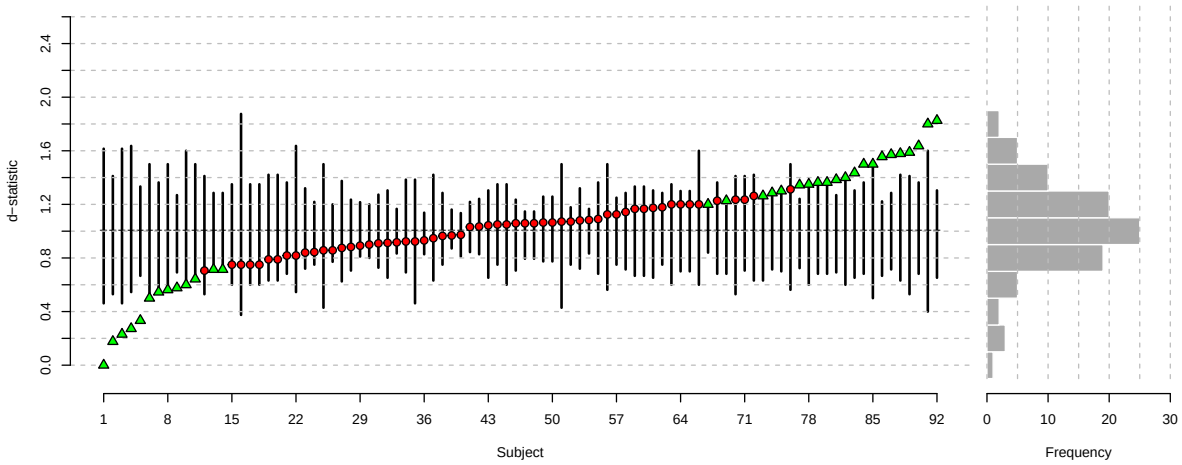


*Figure D.2: Information demand*

*Notes*: The figure displays the distribution of the $d$-statistic across subjects in the CHOICE treatment. A $d$-statistic equal to 1 corresponds to a random choice of senders; a $d$-statistic higher (lower) than 1 corresponds to the selection of informative (uninformative) senders. The left panel ranks the individual $d$-statistics by size (dots and triangles indicate the observed $d$-statistic) and shows a 90% confidence intervals for the subject's $d$-statistic under the null hypothesis of random sender choice (black bars). When the observed $d$-statistic lies outside the black bars (which is indicated with a triangle), one can reject the null hypothesis of a random selection of sender at the 10% confidence level. The right panel shows the overall distribution of the observed $d$-statistic.

right hand side of the left panel), out of a total of 92. The percentage of avoiders in the set of clearly classifiable decision makers is 38.2%, not too far from the 43.5% we found above.

Overall, this analysis confirms that a sizable minority of subjects appears to either avoid informative senders or at 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 $\rho = -0.202$, $p = 0.054$).

## 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 senders 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 sender has been selected.

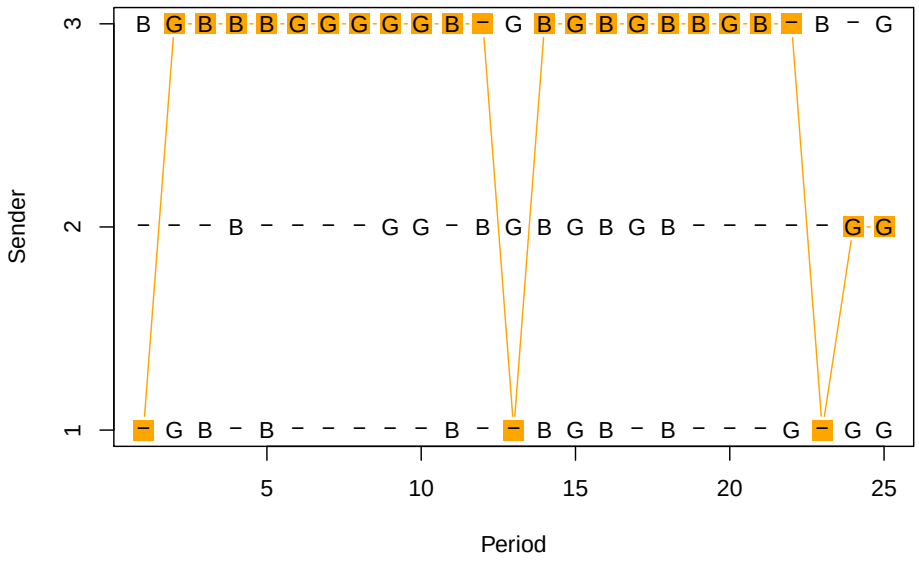


*Figure D.3: Demand for information of a decision maker with d-statistic = 1.23*

*Note*: Figure D.3 illustrates the case of an information seeker who sanctions any transmission of no news ($d$-statistic = 1.23).

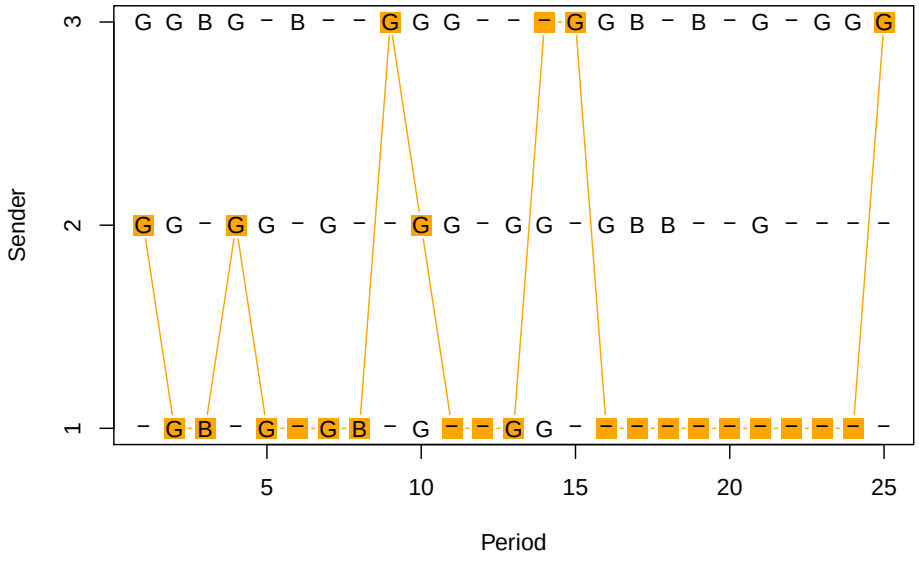


*Figure D.4: Demand for information of a decision maker with d-statistic = 0.75*

*Notes*: Figure D.4 illustrates also the case of an information avoider who is less able to establish a stable relationship with a sender ($d$-statistic = 0.75).

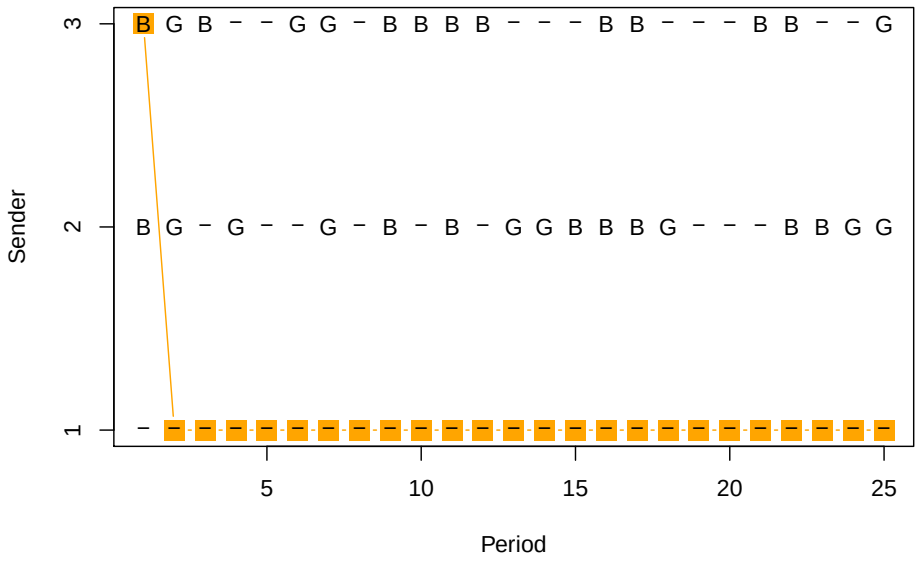


*Figure D.5: Demand for information of a decision maker with d-statistic = 0.18*

*Notes*: Figure D.5 illustrates the case of an information avoider who establishes a long term relationship with a sender who never provides news ($d$-statistic = 0.18).

## Decision to switch sender — Linear probability models

Here we present a third way to analyze demand, namely by the decision to switch to another sender conditional on the information received. This decision is informative of decision makers' information seeking or avoiding strategy. It provides an alternative way to look at information demand based on sender history. Figure D.6 shows the fraction of decision makers that change sender in part 3 in the CHOICE treatment after sender reported good, bad or no news. The left panel shows aggregate results, which demonstrates that switching rates are substantial and vary between 47.3% after bad news, 43.1% after good news, and 56.9% after no news. On aggregate, switching is highest after no news, in line with the idea that most people are information seekers.

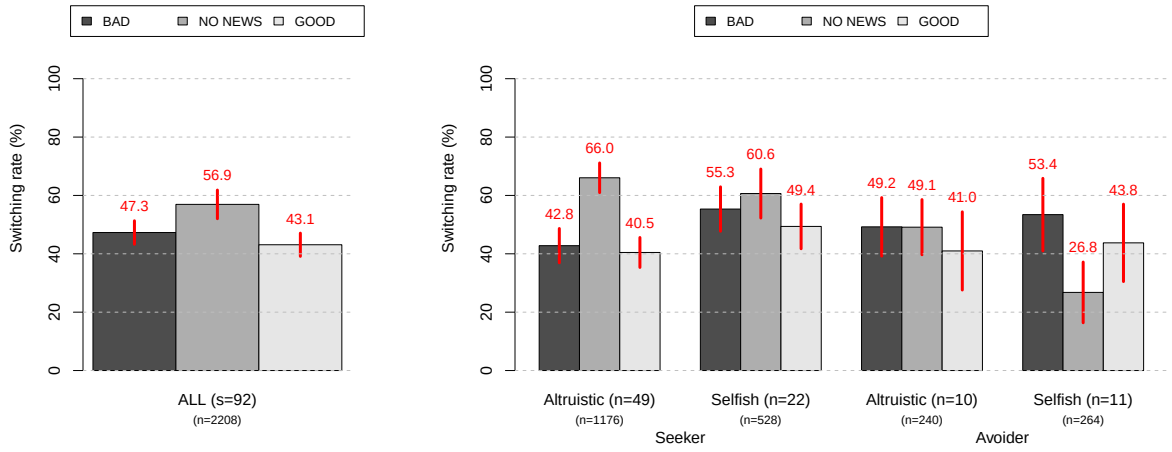


*Figure D.6: Switching senders*

*Notes*: The figure displays the fraction of decision makers that change sender after sender reported good, bad or no news in the CHOICE 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 D.2.

The right panel shows a split by the "type" of decision makers, based on their revealed preferences in part 1. This 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 66.0% after receiving no news, 42.8% after bad news and 40.5% after good news. By contrast, there is a smaller group (Selfish - Avoiders) that does the opposite: the switching rate is 26.8% after receiving no news, 53.4% after bad news and 43.8% 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".

Table D.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.[16] The results of Models 1 and 2 confirm that the baseline category (Altruistic - Seekers) is about 26 percentage points more likely to switch sender 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

[16] In some cases the coefficients do not precisely match the height of the bars in Figure D.6, as the former include demographic control variables, while the latter show pure frequencies.

both by the Selfish and Altruistic Avoiders.

*Table D.2: Decision to switch senders*

| | Model 1<br>Est. (S.E.) | Model 2<br>Est. (S.E.) | Model 3<br>Est. (S.E.) | Model 4<br>Est. (S.E.) |
|---|---|---|---|---|
| (Intercept) | 0.409 (0.052)*** | 0.557 (0.062)*** | 0.405 (0.051)*** | 0.552 (0.062)*** |
| 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) |
| *N 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 sender in part 3 of the CHOICE treatment in each of the 25 periods (24 switching decisions per decision maker). Robust standard errors clustered at group level are in parentheses. *d* for dummy variables. Control variables are: age of the participant (demeaned); gender dummy d(Male); high school grade at the Baccalaureat (BAC) (demeaned); number of past participations in experiments. Period dummies are included with period 2 as the reference category. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; $^\circ \leq 0.1$.

# E Appendix: Additional Analysis of Assortative Matching and Unethical Outcomes

Table E.1 shows the results of linear probability models where the dependent variable is equal to one if a decision maker matches with an Avoider sender. 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 an Avoider increases the probability of matching with an Avoider-sender by 21 percentage points, which is statistically significant at the 5% level. Being a Selfish decision maker does not increase the probability further, and in fact slightly decreases it. This pattern makes sense, as we have shown above that being an Avoider predicts both the demand for ignorance as the supply. In addition, we add an interaction between the two dimensions of the decision-makers' type. This shows that the tendency for assortative matching is somewhat larger for Selfish Avoiders. Overall, this category is almost 30 percentage points more likely to match with an Avoider Sender than the baseline category (Altruistic Seeker), which is statistically significant at the 5% level (see Table Notes).

*Table E.1: Matching of types in the CHOICE treatment*

| | Model 1<br>Est. (S.E.) | Model 2<br>Est. (S.E.) |
|---|---|---|
| (Intercept) | 0.524 (0.103)*** | 0.549 (0.101)*** |
| Selfish DM | -0.059 (0.054) | -0.096 (0.061) |
| Avoider DM | 0.210 (0.105)* | 0.148 (0.111) |
| Selfish DM × Avoider DM | — — | 0.148 (0.132) |
| $Age - \overline{Age}$ | -0.008 (0.002)*** | -0.009 (0.002)*** |
| *d(Male)* | -0.116 (0.049)* | -0.119 (0.047)* |
| $BAC - \overline{BAC}$ | -0.017 (0.011) | -0.018 (0.011) |
| *# of past participations in exp.* | 0.031 (0.015)* | 0.025 (0.015)° |
| *Period dummies* | YES | YES |
| Number of observations | 1150 | 1150 |
| Number of clusters | 12 | 12 |

*Notes*: The regressions are based on linear probability models. The binary dependent variable is the decision maker's choice of a sender who is an Avoider in Part 3 (in each of the 25 rounds). Regressions include only data from the groups where there was at least one sender per type. Robust standard errors clustered at group level are in parentheses. DM for decision maker; *d* for dummy variables. Control variables are: age of the participant (demeaned); gender dummy d(Male); high school grade at the Baccalaureat (BAC) (demeaned); number of past participations in experiments. Period dummies are included with period 1 as the reference category. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; ° $\leq 0.1$. Model 2 Wald test: $\beta = 0.148 + 0.148 = 0.296$, $p = 0.027$).

# F Appendix: Motives behind information supply and demand (questionnaire responses)

In the final questionnaire, we asked both senders and decision makers about their motives in a free form text format. To categorize these motives, we hired three independent raters, who categorized each response into several pre-specified motives. We classify a response if two out of three raters agree on the category. Below we provide more information about the categories.

## F.1 Decision maker motives

At the end of the sessions, decision makers were asked to answer the following question: "According to which principle(s) did you select the sender?". We classified the different comments left by the Senders. Below we illustrate each category by examples of Senders' comments.

**Categories.** Table F.1 summarizes our categories and the instructions we provided to raters.

| Motive | Description | Examples |
|---|---|---|
| Increase my payoff | Mention willingness to increasing own payoffs | "Maximized my payoffs", "I was indifferent, as I always made the decision that paid me 15" |
| Senders' equality | Mentions equality of payoffs between the senders | "I selected in a fair way: in each part I chose a different sender" |
| Information seeking | Mentions the willingness to be informed on the consequences of one's choice on the charity | "I wanted to have information on GD", "the sender who was able to provide information each time", "the frequency of information on the beneficiary" |
| Information avoidance | Mentions the willingness to ignore the consequences of one's choice on the charity | "I did not want to know the consequences", "I only wanted to see animals pictures" |
| Random (No rule) | Mentions randomness in the selection of the sender | "I chose at random" |
| Sender's morality (Honesty) | Mentions a choice based on the honesty or transparency of the sender | "I rewarded transparency" |
| Other | None of the above | - |

*Table F.1: Response categories for decision makers. These motives, descriptions and examples were shared with raters.*

**Results.** Table F.2 summarizes the result of the exercise for both CAP and NoCAP CHOICE treatments.

| | CAP | NoCAP |
|---|---|---|
| Increase my payoff | 1 | 2 |
| Senders' equality | 14 | 11 |
| Information seeking | 43 | 38 |
| Information avoidance | 7 | 10 |
| Random (No rule) | 11 | 12 |
| Sender's morality (Honesty) | 4 | 1 |
| Other (None of the above) | 11 | 16 |
| No coders' agreement | 1 | 2 |
| No explanation provided (blank) | 0 | 0 |
| Total | 92 | 92 |

*Table F.2: Decision makers' classification based of the provided explanation of their reasons to choose a Sender in the CHOICE treatments (classification by majority out of 3 coders: 86.4% of the classifications reached by unanimity, 12.0% by majority, and 1.6% no agreement)*

## F.2 Sender motives

At the end of the sessions, Senders were asked to answer the following question: "According to which principle(s) did you decide to report or not the consequences to the receivers?". We classified the different comments left by the Senders. Below we illustrate each category by examples of Senders' comments.

**Categories.** Table F.3 summarizes our categories and the instructions we provided to raters.

| Motive | Description | Examples |
|---|---|---|
| Strategic | Mention a wish to be selected, increasing own payoffs | "Maximized my payoffs", "I would be more likely to be selected if giving info" |
| Paternalistic | Mentions wish for specific action of decision maker | "I gave information hoping that the decision maker would not cancel the donation", "I counted on them making a donation" |
| Autonomy | Mentions agency or autonomy of the decision maker | "I gave information to let them make their own decision", "let them decide according to their values", "giving them responsibility" |
| Help the decision maker | Mentions wish to help the decision maker in some way or fulfill preferences of decision maker | "Make their decisions easy", "Relieve their conscience", "they would prefer remaining ignorant/having information", "I gave information according to their earlier choices" |
| Rule based | Mentions general (moral) rule | "Transparency is important", "Always give information" |
| No Rule | Mentions randomness in the decision to send the picture of a beneficiary or the picture of a cute animal | "I chose at random" |

*Table F.3: Response categories for senders. These motives, descriptions and examples were shared with raters.*

**Results.** Table F.4 summarizes the result of the exercise split by CHOICE and RANDOM and CAP and NoCAP. Table F.5 aggregates these results over the CAP and NoCAP treatments, and hence focuses on CHOICE versus RANDOM.

| | Choice - CAP | Choice - NoCAP | Random - CAP | Random - NoCAP |
|---|---|---|---|---|
| Strategic | 4 | 8 | 1 | 1 |
| Make DM decision easy (Help DM) | 12 | 8 | 2 | 6 |
| Paternalistic | 6 | 7 | 13 | 15 |
| Autonomy of DM | 5 | 6 | 6 | 15 |
| Moral rule | 14 | 11 | 13 | 8 |
| Random (No rule) | 10 | 12 | 11 | 7 |
| Other (None of the above) | 11 | 13 | 21 | 14 |
| No coders' agreement | 5 | 3 | 2 | 0 |
| No explanation provided (blank) | 2 | 1 | 0 | 0 |
| Total | 69 | 69 | 69 | 66 |

*Table F.4: Senders' classification based of the provided explanation of their reasons to share (or not) information (classification by majority out of three coders: 64.8% of the classifications reached by unanimity, 31.5% by majority, and 5.7% no agreement)*

| | Choice | Random |
|---|---|---|
| Strategic | 12 | 2 |
| Make DM decision easy (Help DM) | 20 | 8 |
| Paternalistic | 13 | 28 |
| Autonomy of DM | 11 | 21 |
| Moral rule | 25 | 21 |
| Random (No rule) | 22 | 18 |
| Other (None of the above) | 24 | 35 |
| No coders' agreement | 8 | 2 |
| No explanation provided (blank) | 3 | 0 |
| Total | 138 | 135 |

*Table F.5: Senders' classification based of the provided explanation of their reasons to share (or not) information (classification by majority out of three coders: 64.8% of the classifications reached by unanimity, 31.5% by majority, and 5.7% no agreement)*

# G Appendix: Analyses of the NoCAP treatments (without cute animal pictures)

## G.1 Comparison of the CAP and NoCAP treatments

*Table G.1: Unethical outcomes and ignorance consumption in CAP and NoCAP*

| | Unethical outcomes | | Ignorance consumption | |
|---|---|---|---|---|
| | Model 1 | Model 2 | Model 3 | Model 4 |
| | Est. (S.E.) | Est. (S.E.) | Est. (S.E.) | Est. (S.E.) |
| Constant | 0.840 (0.047)*** | 0.851 (0.052)*** | 0.374 (0.042)*** | 0.379 (0.050)*** |
| d(NoCAP) | 0.025 (0.028) | 0.019 (0.037) | -0.035 (0.029) | -0.046 (0.043) |
| d(CHOICE) | | -0.021 (0.046) | | -0.014 (0.047) |
| d(NoCAP) × d(CHOICE) | | 0.013 (0.057) | | 0.022 (0.058) |
| Altruistic DM | -0.432 (0.026)*** | -0.433 (0.026)*** | -0.025 (0.020) | -0.024 (0.021) |
| Avoider DM | 0.100 (0.035)** | 0.101 (0.035)** | 0.074 (0.030)* | 0.074 (0.031)* |
| $Age - \overline{Age}$ | -0.002 (0.002) | -0.002 (0.002) | -0.002 (0.002) | -0.002 (0.002) |
| *d(Male)* | 0.026 (0.028) | 0.026 (0.028) | 0.005 (0.019) | 0.006 (0.019) |
| $BAC - \overline{BAC}$ | 0.005 (0.007) | 0.005 (0.007) | -0.002 (0.005) | -0.002 (0.005) |
| *# of past participations in exp.* | -0.002 (0.009) | -0.002 (0.009) | -0.007 (0.008) | -0.007 (0.008) |
| *Period dummies* | YES | YES | YES | YES |
| Number of observations | 4498 | 4498 | 8875 | 8875 |
| Number of clusters | 91 | 91 | 91 | 91 |

*Notes*: The regressions are based on linear probability models. The binary dependent variable in Model 1 and 2 is the cancellation of the donation in Part 3 in each of the 25 rounds. These models include only data where the state is bad. The binary dependent variable in Model 3 and 4 takes value one when the decision maker is observing NO NEWS. Robust standard errors clustered at group level are in parentheses. DM for decision maker; *d* for dummy variables. Control variables are: age of the participant (demeaned); gender dummy d(Male); high school grade at the Baccalaureat (BAC) (demeaned); number of past participations in experiments. Period dummies are included with period 1 as the reference category. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; $^{\circ} \leq 0.1$.

*Table G.2: Suppression of bad and good news in CAP and NoCAP treatments*

| | Bad news<br>Model 1<br>Est. (S.E.) | Good news<br>Model 2<br>Est. (S.E.) |
|---|---|---|
| (Intercept) | 0.177 (0.053)*** | 0.241 (0.048)*** |
| d(NoCAP) | -0.046 (0.037) | -0.065 (0.032)* |
| Altruistic | -0.076 (0.033)* | -0.077 (0.032)* |
| Avoider | 0.146 (0.051)** | 0.097 (0.049)* |
| Belief # ignorant | 0.015 (0.004)*** | 0.006 (0.004) |
| $Age - \overline{Age}$ | -0.003 (0.002) | 0.0002 (0.003) |
| *d(Male)* | 0.005 (0.030) | 0.017 (0.033) |
| $BAC - \overline{BAC}$ | -0.006 (0.007) | 0.003 (0.006) |
| *# of past participations in exp.* | -0.012 (0.012) | -0.008 (0.011) |
| *Period dummies* | YES | YES |
| Number of observations | 10731 | 10438 |
| Number of clusters | 91 | 91 |

*Notes*: The regressions are based on linear probability models. The binary dependent variable in Model 1 is the senders choice to suppress bad news in Part 3 (in each of the 25 rounds), the binary dependent variable in Model 2 is the senders choice to suppress good news in Part 3 (in each of the 25 rounds). Robust standard errors clustered at group level are in parentheses. DM for decision maker; *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. Control variables are: age of the participant (demeaned); gender dummy d(Male); high school grade at the Baccalaureat (BAC) (demeaned); number of past participations in experiments. Period dummies are included with period 1 as the reference category. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; $^{\circ} \leq 0.1$.

## G.2 Replication of the analyses made for the CAP treatment

*Table G.3: Unethical outcomes and ignorance consumption (NoCAP version of Table 2)*

| | Unethical outcomes | | Ignorance consumption | |
|---|---|---|---|---|
| | Model 1 | Model 2 | Model 3 | Model 4 |
| | Est. (S.E.) | Est. (S.E.) | Est. (S.E.) | Est. (S.E.) |
| (Intercept) | 0.858 (0.054)*** | 0.827 (0.052)*** | 0.371 (0.060)*** | 0.355 (0.063)*** |
| d(RANDOM) | 0.003 (0.031) | 0.090 (0.028)** | -0.014 (0.034) | 0.019 (0.047) |
| Avoider DM | 0.130 (0.046)** | 0.181 (0.050)*** | 0.087 (0.052)° | 0.262 (0.073)*** |
| Avoider DM × d(RANDOM) | | -0.078 (0.085) | | -0.333 (0.081)*** |
| Altruistic DM | -0.454 (0.034)*** | -0.392 (0.052)*** | -0.021 (0.027) | -0.016 (0.042) |
| Altruistic DM × d(RANDOM) | | -0.130 (0.060)* | | 0.017 (0.050) |
| $Age - \overline{Age}$ | 0.002 (0.003) | 0.002 (0.003) | 0.003 (0.003) | 0.003 (0.002) |
| *d(Male)* | 0.020 (0.035) | 0.017 (0.037) | -0.005 (0.026) | -0.017 (0.024) |
| $BAC - \overline{BAC}$ | 0.002 (0.010) | 0.002 (0.010) | -0.0005 (0.006) | 0.001 (0.006) |
| *# of past participations in exp.* | -0.011 (0.012) | -0.012 (0.012) | -0.019 (0.012) | -0.021 (0.010)* |
| *Period dummies* | YES | YES | YES | YES |
| Number of observations | 2237 | 2237 | 4450 | 4450 |
| Number of clusters | 45 | 45 | 45 | 45 |

*Notes*: The regressions are based on linear probability models. The binary dependent variable in Model 1 and 2 is the cancellation of the donation in Part 3 in each of the 25 rounds (i.e., the choice of Option 2). These models include only data where the state is bad. The binary dependent variable in Models 3 and 4 takes value one when the decision maker is observing NO NEWS. Robust standard errors clustered at group level are in parentheses. DM for decision maker; *d* for dummy variables. Control variables are: age of the participant (demeaned); gender dummy d(Male); high school grade at the Baccalaureat (BAC) (demeaned); number of past participations in experiments. Period dummies are included with period 1 as the reference category. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; ° $\leq 0.1$.

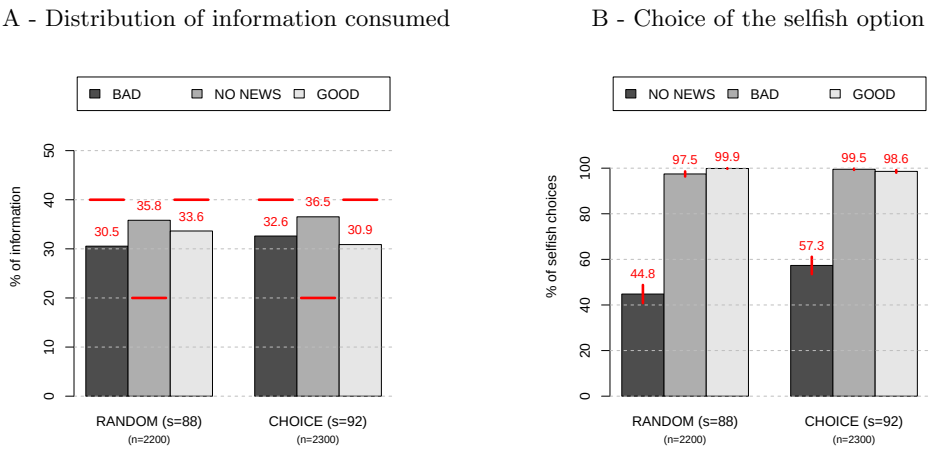


*Figure G.1: Information consumption and choices of the decision makers (NoCAP version of Figure 2)*

*Notes*: Panel A 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 senders. Panel B 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 ($n$).

*Table G.4: Suppression of bad news by senders (NoCAP version of Table 3)*

| | Model 1<br>Est. (S.E.) | Model 2<br>Est. (S.E.) | Model 3<br>Est. (S.E.) | Model 4<br>Est. (S.E.) |
|---|---|---|---|---|
| (Intercept) | 0.069 (0.070) | 0.057 (0.073) | 0.041 (0.072) | 0.061 (0.066) |
| Belief # ignorant | 0.018 (0.005)*** | 0.018 (0.005)*** | 0.021 (0.006)** | 0.015 (0.006)* |
| d(CHOICE) | — — | 0.021 (0.044) | 0.051 (0.048) | 0.040 (0.048) |
| d(CHOICE)×Belief # ignorant | — — | — — | -0.006 (0.010) | -0.006 (0.008) |
| Altruistic | — — | — — | — — | -0.031 (0.041) |
| Avoider | — — | — — | — — | 0.173 (0.071)* |
| $Age - \overline{Age}$ | -0.004 (0.006) | -0.004 (0.006) | -0.005 (0.006) | -0.008 (0.006) |
| *d(Male)* | 0.024 (0.046) | 0.026 (0.046) | 0.026 (0.046) | 0.038 (0.043) |
| $BAC - \overline{BAC}$ | 0.002 (0.009) | 0.001 (0.009) | 0.001 (0.010) | -0.003 (0.010) |
| *# of past participations in exp.* | -0.013 (0.018) | -0.013 (0.018) | -0.013 (0.018) | -0.011 (0.018) |
| *Period dummies* | YES | YES | YES | YES |
| Number of observations | 5342 | 5342 | 5342 | 5342 |
| Number of clusters | 45 | 45 | 45 | 45 |

*Notes*: These regressions are based on linear probability models. The binary dependent variable is the sender's choice to suppress bad news in Part 3 in each of the 25 rounds. 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. Control variables are: age of the participant (demeaned); gender dummy d(Male); high school grade at the Baccalaureat (BAC) (demeaned); number of past participations in experiments. Period dummies are included with period 1 as the reference category. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; ° $\leq 0.1$.

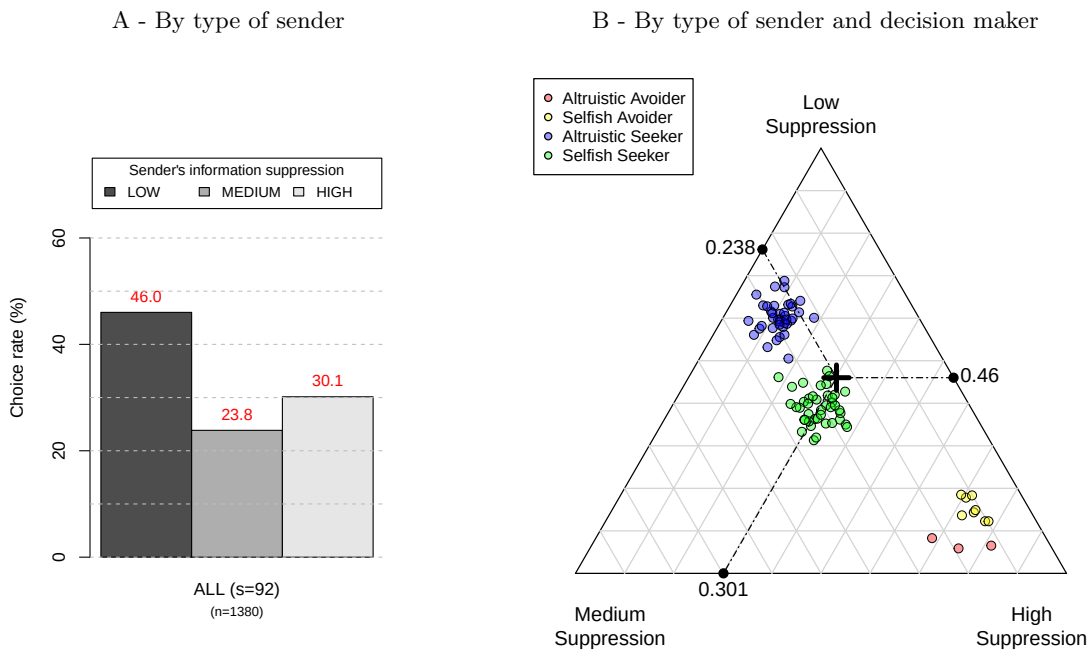


*Figure G.2: Predicted probability to choose a given sender (NoCAP version of Figure 3)*

*Notes*: The figure displays the frequency of choices of the three senders in the CHOICE treatment. Senders are ranked (low, medium, and high) according to the relative level of ignorance they provided in the previous 10 periods. Panel A shows displays the predicted probability to choose the sender providing the highest, intermediate, and lowest level of suppression for each of the 92 decision makers. Predicted probabilities are based on Model 1 of Table G.5 of appendix G. Panel B adds the type of the decision maker. The color of the dot captures the type of the decision maker based on decisions in Part 1. The cross shows the average frequency of choice.

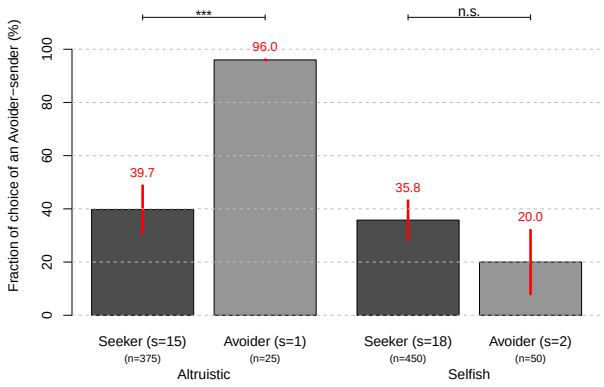


*Figure G.3: Assortative matching - choice of an avoider sender by decision maker type (NoCAP version of Figure 4)*

*Notes*: The figure displays the frequency of choice of an Avoider sender. Decision makers are split by type, as defined based on decisions in Part 1. The frequencies are calculated on the subset of groups where there is at least one and at most two Avoiders among the senders. Vertical lines represent standard errors based on a linear probability model with clustering at the group level. The two factors on the x-axis and their interaction are the only explanatory variables. Pairwise comparisons reported above the bars are based on a Wald test performed using this estimated model. Signific. codes: *** $p < 0.001$; n.s. means $p \geq 0.05$.

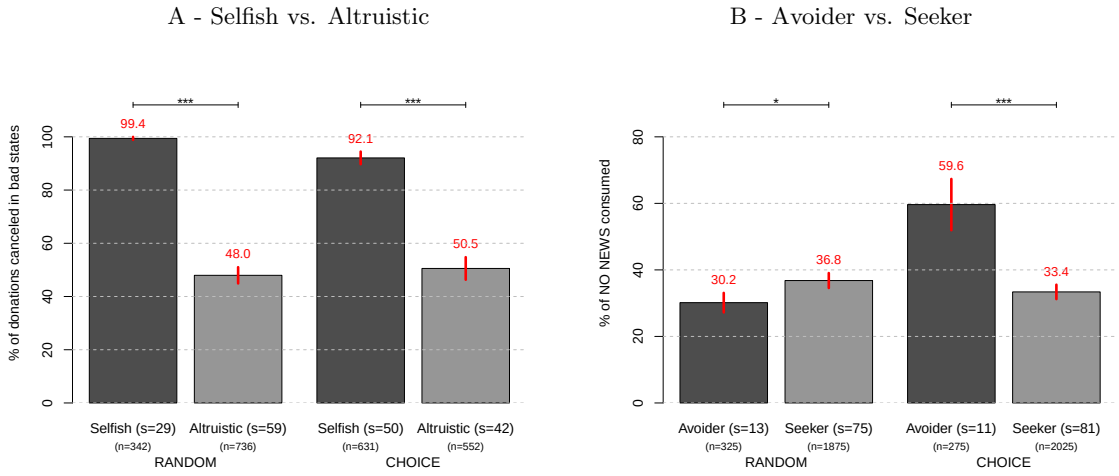


*Figure G.4: Selfish decisions in the bad state and ignorance consumption by condition and decision-maker type (NoCAP version of Figure 5).*

*Notes*: Panel A displays the fraction of selfish decisions in the different experimental conditions. The fraction of selfish decisions in the RANDOM and CHOICE treatments are computed using only the cases where the state is bad. Panel B displays the fraction of NO NEWS consumption in the different experimental conditions. Panel A splits decision makers into Selfish and Altruistic. Panel B splits them into Avoider and Seeker, based on their decision in Part 1. In all panels, vertical lines represent standard errors based on a linear probability model with clustering at the group level. In all models, the two factors on the x-axis and their interaction are the only explanatory variables. Pairwise comparisons reported above the bars are based on a Wald test performed using these estimated models. Signific. codes: * $p \leq 0.05$; ** $p \leq 0.01$; *** $p \leq 0.001$; n.s. means $p > 0.10$.

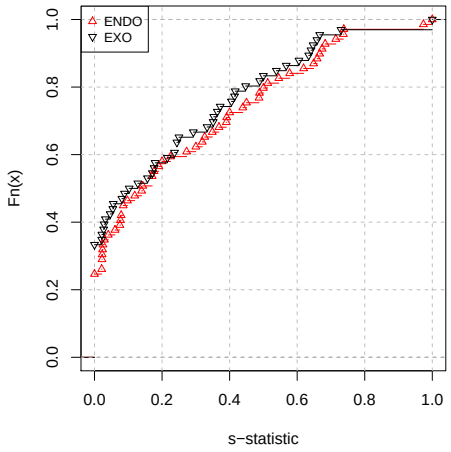


*Figure G.5: Suppression of bad news by senders (NoCAP version of Figure C.1)*

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

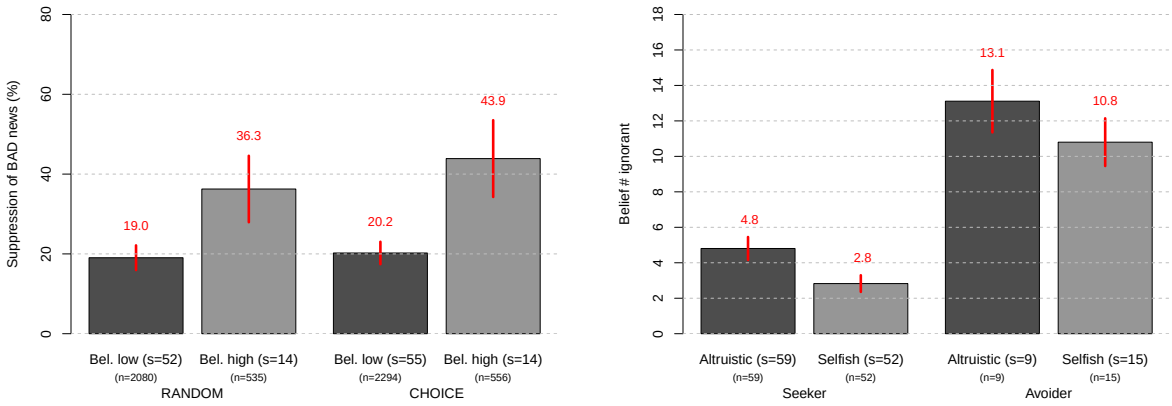


*Figure G.6: Sender beliefs and suppression of bad news (NoCAP version of Figure C.2)*

*Notes*: The left panel shows the impact of sender's beliefs about the number of subjects (in the session) that want to remain uninformed on the rate of suppression of bad news. Belief high (low) is for senders that believe the majority wants to remain ignorant (wants to know). The right panel shows the average belief about the number of subjects that want to remain ignorant conditional on the sender's preferences as revealed in Part 1 of the experiment. Vertical bars show the standard errors. In the left panel, 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 G.4. Compared to the table, the model only includes the treatment dummy, a dummy that is equal to 1 when the sender believes that the majority wants to remain uninformed, and their interaction as explanatory variables.

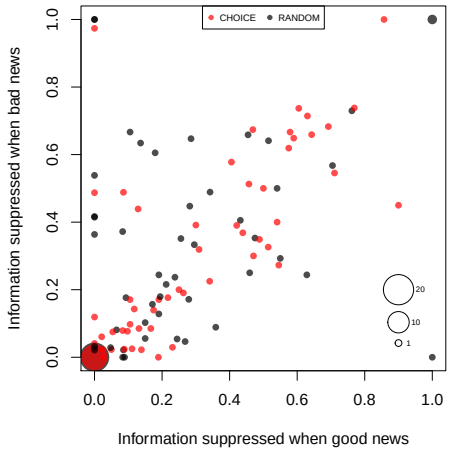


*Figure G.7: Individual suppression rates of good and bad news by treatment (NoCAP version of Figure C.3)*

*Notes*: The figure displays suppression of good news (x-axis) and bad news (y-axis). Each dot is a sender. Multiple senders with the same suppression patterns are indicated by bigger dots.

*Table G.5: Probability to choose the sender that provided the highest, intermediate and lowest level of ignorance (NoCAP version of Table D.1)*

| | Model 1<br>Est. (S.E.) | Model 2<br>Est. (S.E.) |
|---|---|---|
| Medium ignorance | | |
| (intercept) | -1.011 (0.122)*** | -1.932 (0.503)*** |
| Selfish (DM) | 0.616 (0.163)*** | 0.105 (0.203) |
| Avoider (DM) | 1.819 (0.903)* | 0.151 (1.051) |
| Selfish (DM) × Avoider (DM) | -1.960 (0.968)* | 0.517 (1.133) |
| $\sigma_M$ | 1.186 (0.127)*** | 0.915 (0.137)*** |
| High ignorance | | |
| (intercept) | -2.818 (0.244)*** | -2.319 (0.728)** |
| Selfish (DM) | 2.172 (0.239)*** | 1.973 (0.312)*** |
| Avoider (DM) | 6.089 (0.805)*** | 5.247 (1.003)*** |
| Selfish (DM) × Avoider (DM) | -3.253 (0.832)*** | -3.653 (1.098)*** |
| $\sigma_H$ | 1.962 (0.159)*** | 2.598 (0.268)*** |
| *Group dummies* | NO | YES |
| Log-Likelihood | -1197.3 | -1161.7 |
| 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 senders ordered by the amount of ignorance supplied in the previous 10 periods (the baseline alternative is the sender 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. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; ° $\leq 0.1$.

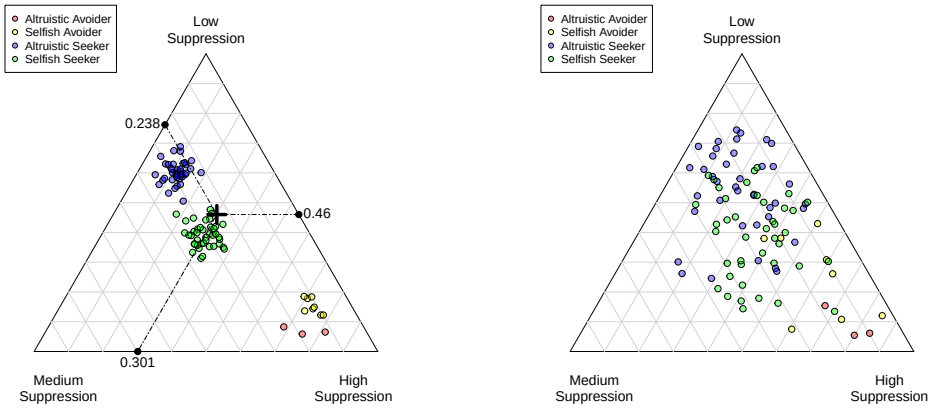


*Figure G.8: Predicted probability of the decision makers' choice of sender (NoCAP version of Figure D.1)*

*Notes*: The figures display the predicted probability to choose the sender 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 G.5 and the figure on the right reports predictions based on Model 2 of Table G.5. The color of the dots captures the type of the decision maker elicited in Part 1. The cross shows the average frequency of choice.

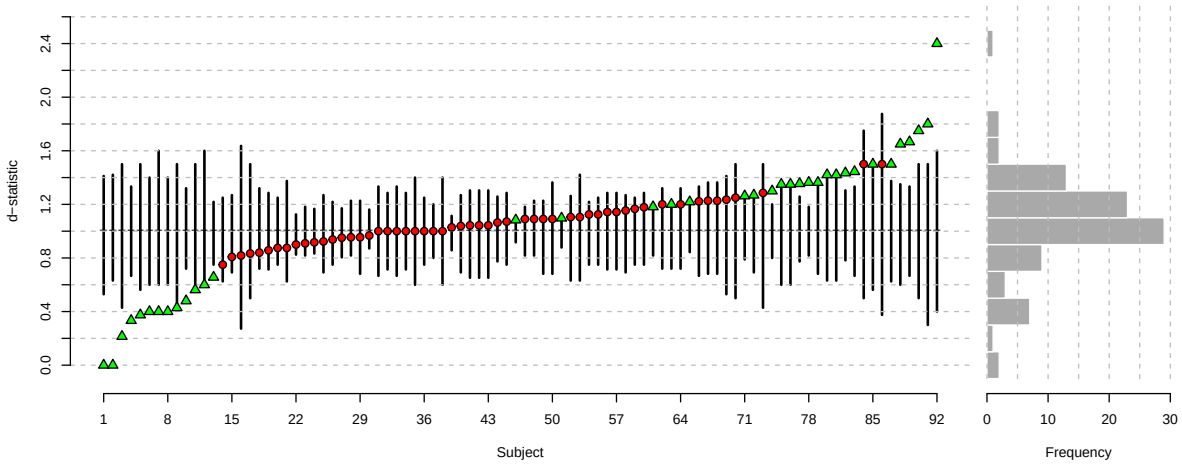


*Figure G.9: Information demand (NoCAP version of Figure D.2)*

*Notes*: The figure displays the distribution of the $d$-statistic across subjects in the CHOICE treatment. A $d$-statistic equal to 1 corresponds to a random choice of senders; a $d$-statistic higher (lower) than 1 corresponds to the selection of informative (uninformative) senders. The left panel ranks the individual $d$-statistics by size (dots and triangles indicate the observed $d$-statistic) and shows a 90% confidence intervals for the subject's $d$-statistic under the null hypothesis of random sender choice (black bars). When the observed $d$-statistic lies outside the black bars (which is indicated with a triangle), one can reject the null hypothesis of a random selection of sender at the 10% confidence level. The right panel shows the overall distribution of the observed $d$-statistic.

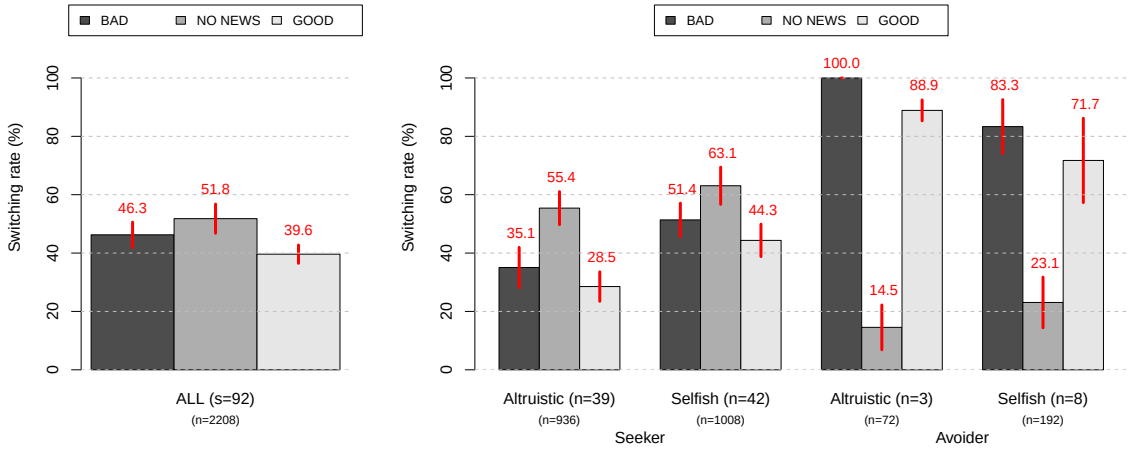


*Figure G.10: Switching senders (NoCAP version of Figure D.6)*

*Notes*: The figure displays the fraction of decision makers that change sender after sender reported good, bad or no news in the CHOICE 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 G.6.

*Table G.6: Decision to switch senders (NoCAP version of Table D.2)*

| | Model 1<br>Est. (S.E.) | Model 2<br>Est. (S.E.) | Model 3<br>Est. (S.E.) | Model 4<br>Est. (S.E.) |
|---|---|---|---|---|
| (Intercept) | 0.293 (0.050)*** | 0.445 (0.097)*** | 0.285 (0.051)*** | 0.438 (0.096)*** |
| GOOD news × Selfish | 0.143 (0.076)° | 0.130 (0.072)° | 0.158 (0.082)° | 0.140 (0.076)° |
| GOOD news × Avoider | 0.333 (0.152)* | 0.330 (0.134)* | 0.604 (0.064)*** | 0.496 (0.064)*** |
| GOOD news × Selfish × Avoider | — — | — — | -0.330 (0.200)° | -0.200 (0.178) |
| BAD news | 0.064 (0.045) | 0.069 (0.049) | 0.065 (0.046) | 0.070 (0.051) |
| BAD news × Selfish | 0.150 (0.087)° | 0.125 (0.079) | 0.163 (0.089)° | 0.133 (0.082) |
| BAD news × Avoider | 0.377 (0.088)*** | 0.371 (0.099)*** | 0.649 (0.069)*** | 0.515 (0.079)*** |
| BAD news × Selfish × Avoider | — — | — — | -0.330 (0.120)** | -0.174 (0.139) |
| No news | 0.260 (0.044)*** | 0.257 (0.039)*** | 0.269 (0.046)*** | 0.260 (0.042)*** |
| No news × Selfish | 0.078 (0.070) | 0.087 (0.070) | 0.077 (0.083) | 0.091 (0.083) |
| No news × Avoider | -0.403 (0.081)*** | -0.384 (0.083)*** | -0.409 (0.094)*** | -0.369 (0.083)*** |
| No news × Selfish × Avoider | — — | — — | 0.009 (0.144) | -0.023 (0.140) |
| $Age - \overline{Age}$ | — — | 0.005 (0.035) | — — | 0.005 (0.035) |
| *d(Male)* | — — | -0.077 (0.086) | — — | -0.075 (0.086) |
| $BAC - \overline{BAC}$ | — — | -0.007 (0.021) | — — | -0.007 (0.021) |
| *# of past participations in exp.* | — — | 0.018 (0.031) | — — | 0.018 (0.031) |
| *Period dummies* | NO | YES | NO | YES |
| Number of observations | 2208 | 2160 | 2208 | 2160 |
| 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 sender in part 3 of the CHOICE treatment in each of the 25 periods (24 switching decisions per decision maker). Robust standard errors clustered at group level are in parentheses. *d* for dummy variables. Control variables are: age of the participant (demeaned); gender dummy d(Male); high school grade at the Baccalaureat (BAC) (demeaned); number of past participations in experiments. Period dummies are included with period 2 as the reference category. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; ° $\leq 0.1$.

*Table G.7: Matching of types in the CHOICE treatment (NoCAP version of Table E.1)*

| | Model 1<br>Est. (S.E.) |
|---|---|
| (Intercept) | 0.515 (0.125)*** |
| Selfish DM | -0.133 (0.109) |
| Avoider DM | 0.037 (0.248) |
| Selfish DM × Avoider DM | — — |
| $Age - \overline{Age}$ | 0.067 (0.023)** |
| *d(Male)* | 0.088 (0.106) |
| $BAC - \overline{BAC}$ | 0.053 (0.028)° |
| *# of past participations in exp.* | -0.010 (0.050) |
| *Period dummies* | YES |
| Number of observations | 875 |
| Number of clusters | 9 |

*Notes*: The regression is based on a linear probability model. The binary dependent variable is the decision maker's choice of a sender who is an Avoider in Part 3 (in each of the 25 rounds). The regression includes only data from the groups where there was at least one sender per type. Robust standard errors clustered at group level are in parentheses. DM for decision maker; *d* for dummy variables. Control variables are: age of the participant (demeaned); gender dummy d(Male); high school grade at the Baccalaureat (BAC) (demeaned); number of past participations in experiments. Period dummies are included with period 1 as the reference category. *** $\leq 0.001$; ** $\leq 0.01$; * $\leq 0.05$; ° $\leq 0.1$.