



HAL
open science

The Social Construction of Ignorance: Experimental Evidence

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

► **To cite this version:**

Ivan Soraperra, Joël van der Weele, Marie Claire Villeval, Shaul Shalvi. The Social Construction of Ignorance: Experimental Evidence. *Games and Economic Behavior*, 2023, 138, pp.197-213. 10.1016/j.geb.2022.12.002 . hal-03899658

HAL Id: hal-03899658

<https://hal.science/hal-03899658>

Submitted on 15 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Social Construction of Ignorance: Experimental Evidence*

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

December 13, 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. Supporting materials like the Online Appendix, data, analysis scripts, IRB approval and earlier versions of the paper can be found at <https://osf.io/3r62f/>. 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. Joël 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, University 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.

¹A 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-

²A 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](#)).

³There 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 Online Appendix A. Further supporting materials like data, analysis scripts, IRB approval and earlier versions of the paper can be found at <https://osf.io/3r62f/>.

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.

⁴We 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.

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 instructions in Online 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”.

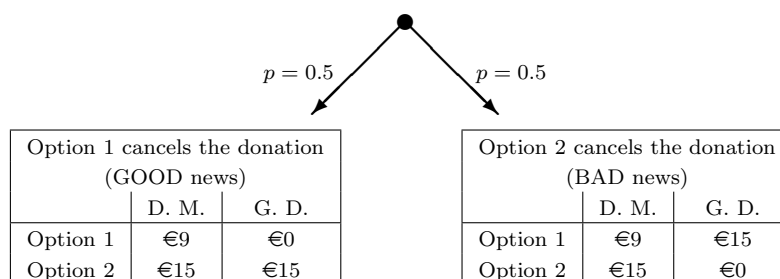


Figure 1: Summary of the decision maker’s and charity’s payoffs

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 Online 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 beginning 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 Online Appendix 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.

⁵We thank two anonymous referees for suggesting these treatments.

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 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 Online 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 Online 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

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

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

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; 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

⁸Note 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>.

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' 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.⁹

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

⁹An earlier version of the paper featured a different version of this conjecture, which stated that CHOICE would reduce overall ethical behavior by making it easier to choose uninformative senders. We thank a referee for pointing out that CHOICE will also raise the possibility to choose informative senders, and hence its impact on ethical behavior is ex-ante unclear. Our current conjecture takes into account these diverging effects.

Table 1: Distribution of types by treatment

Type	(CAP)			(NoCAP)	
	Cute animal picture			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 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.

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.

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

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) ^o	-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)
<i>Age</i> - \overline{Age}	-0.005 (0.002)	-0.005 (0.002)*	-0.004 (0.002)*	-0.004 (0.001)**
<i>d(Male)</i>	-0.001 (0.033)	0.006 (0.031)	-0.007 (0.026)	-0.008 (0.025)
<i>BAC</i> - \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)
<i>Period dummies</i>	YES	YES	YES	YES
Number of observations	3232	3232	6300	6300
Number of clusters	120	120	120	120

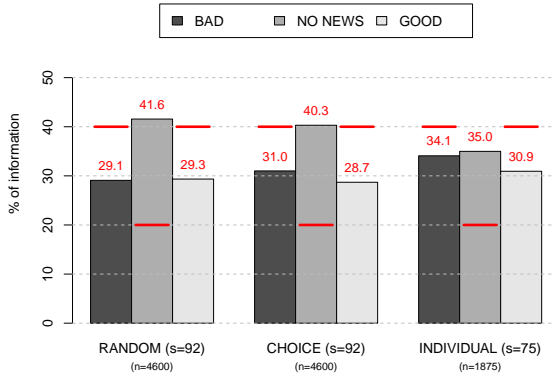
Notes: The regressions are based on linear probability models. The binary dependent variable in Model 1 and 2 is the cancellation of the donation in Part 3 in each of the 25 rounds (i.e., the choice of Option 2). These models include only data where the state is bad. The binary dependent variable in 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. *** ≤ 0.001 ; ** ≤ 0.01 ; * ≤ 0.05 ; ^o ≤ 0.1 .

20% no news.¹⁰ 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).¹¹ 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 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.

¹⁰The 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.

¹¹Repeating the test using the realized frequencies instead of the theoretical ones gives the same results.

A - Distribution of information consumed



B - Choice of the selfish option

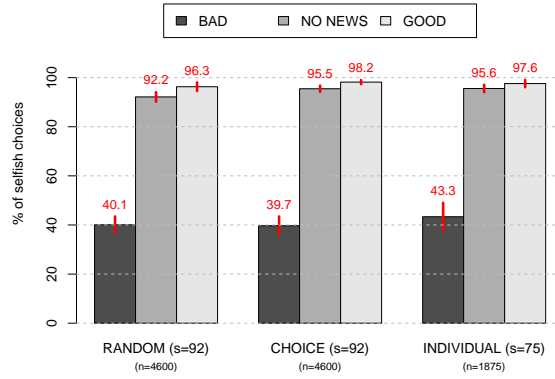


Figure 2: Information consumption and choices of the decision makers

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).

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.

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.¹² To measure suppression of bad news, we look at the fraction of bad states observed

¹²Although 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

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 sender suppression. Figure C.1 in Online 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.¹³ 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 in the Online Appendix G). Alternatively,

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. We provide more detail in Online Appendix C, where Figure C.3 shows that symmetric suppression of news is quite common. 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.

¹³There 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 Kolmogorov-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.

Table 3: Suppression of bad news by senders

	Model 1	Model 2	Model 3	Model 4
	Est. (S.E.)	Est. (S.E.)	Est. (S.E.)	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)*
<i>Age</i> – <i>Age</i>	-0.003 (0.003)	-0.003 (0.003)	-0.003 (0.003)	-0.001 (0.003)
<i>d(Male)</i>	0.003 (0.046)	0.003 (0.045)	-0.004 (0.043)	-0.043 (0.042)
<i>BAC</i> – <i>BAC</i>	-0.009 (0.011)	-0.009 (0.011)	-0.010 (0.011)	-0.013 (0.010)
<i># of past participations in exp.</i>	-0.011 (0.017)	-0.011 (0.017)	-0.009 (0.017)	-0.006 (0.017)
<i>Period dummies</i>	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. *** ≤ 0.001 ; ** ≤ 0.01 ; * ≤ 0.05 ; [◦] ≤ 0.1 .

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.¹⁴

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 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 treat-

¹⁴In 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.

ments 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 Online Appendix 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.¹⁵ 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

¹⁵As 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.

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 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.

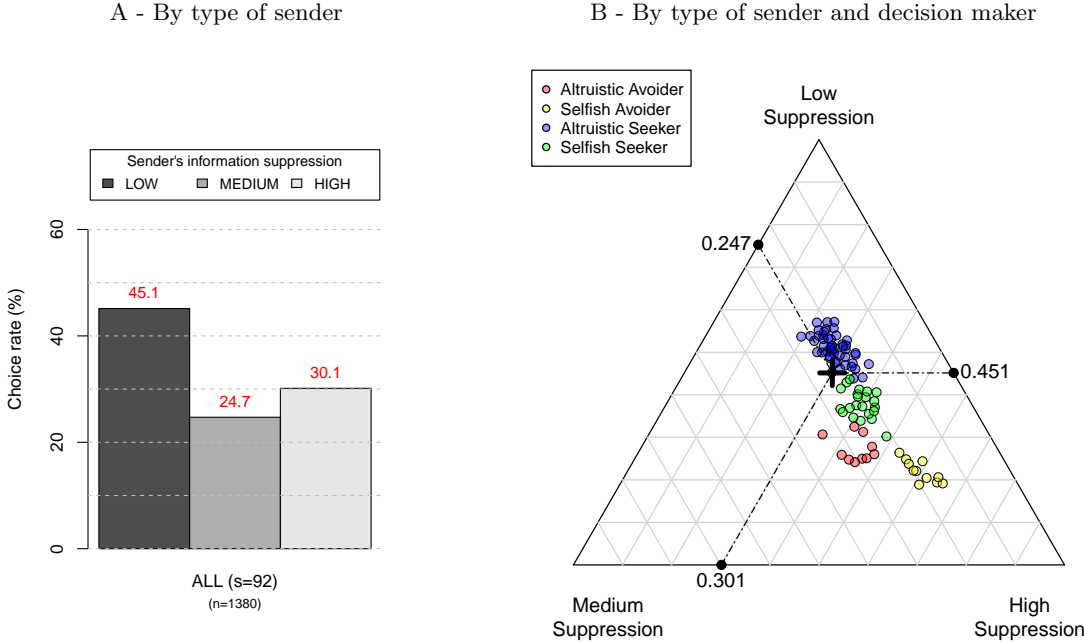


Figure 3: Predicted probability to choose a given sender

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 Online 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 disaggregates 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

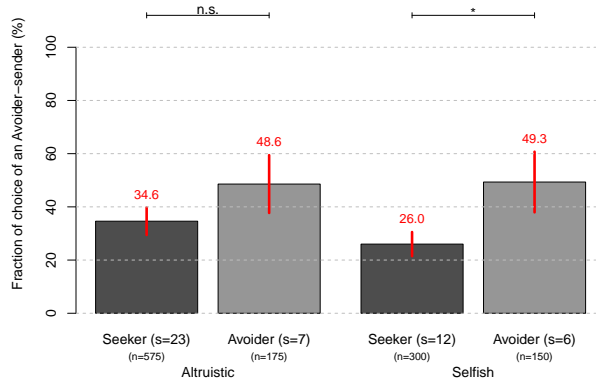


Figure 4: Assortative matching - choice of an avoyer sender by decision maker type.

Notes: The figure displays the frequency of choice of an Avoyer 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 Avoyers 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 $p > 0.05$.

results reported in models (1) and (2) of Table E.1 in Appendix E, is coherent with the fact that being an Avoyer 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, Avoyer-type decision makers match significantly more often with Avoyer type senders.*

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.

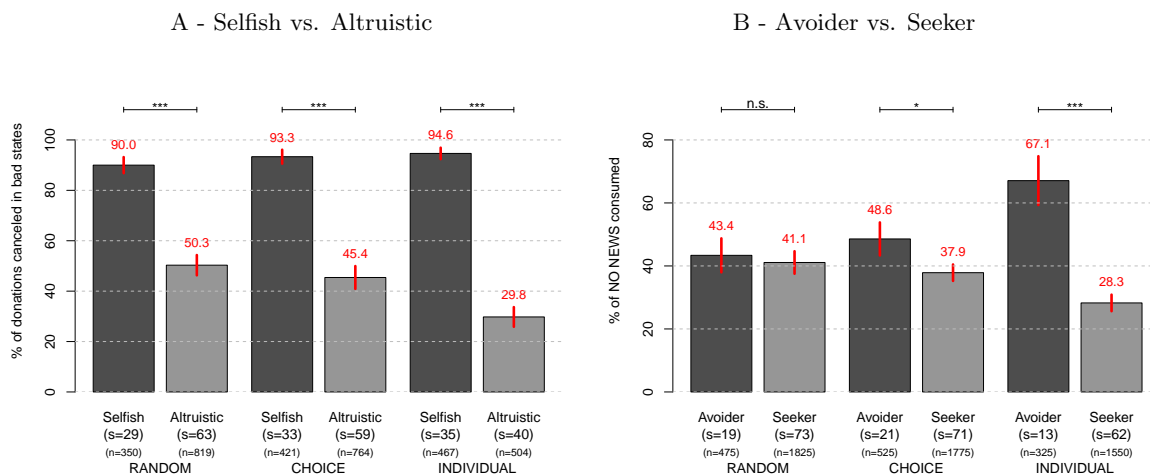


Figure 5: Selfish decisions in the bad state and ignorance consumption by condition and decision-maker type.

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 $p > 0.10$.

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 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 comparison 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.*

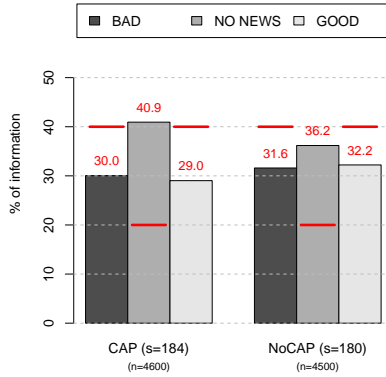
5 The role of cute animal pictures

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.

5.1 Aggregate results

Online 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 in Online Appendix G), 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%.

A - Distribution of information consumed



B - Choice of the selfish option

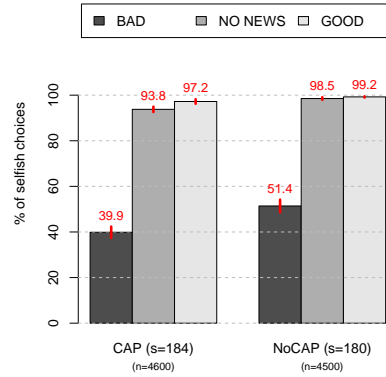


Figure 6: Information consumption and choices of the decision makers in the CAP and NoCAP treatments

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).

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 in Online Appendix G (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

Online Appendix G shows the NoCAP equivalents of all figures and tables presented in the above analysis 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 ex-

periment. 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 Online Appendix 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 Online Appendix 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 Online 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.
- ISMAILOV, 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. CHARNNESS (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.