




ORIGINAL PAPER

Spoiling the party: Experimental evidence on the willingness to transmit inconvenient ethical information

Jantsje M. Mol^{1,2} , Ivan Soraperra³  and Joël J. van der Weele^{1,2} 

¹Center for Research in Experimental Economics and Political Decision Making (CREED), University of Amsterdam, 1018 WB Amsterdam, Netherlands

²Tinbergen Institute, Gustav Mahlerplein 117, 1082 MS Amsterdam, The Netherlands

³Max Planck Institute for Human Development, Center for Humans and Machines, Lentzeallee 94, 14195 Berlin, Germany

Corresponding author: Jantsje M. Mol; Email: j.m.mol@uva.nl

(Received 5 December 2024; revised 20 February 2025; accepted 27 February 2025)

Abstract

Information about the consequences of our consumption choices can be unwelcome, and people sometimes avoid it. Thus, when people possess information that is inconvenient for *another* person, they may face a dilemma about whether to inform them. We introduce a simple and portable experimental game to analyze the transmission of inconvenient information. In this game, a Sender can, at a small cost, inform a Receiver about a negative externality associated with a tempting and profitable action for the Receiver. The results from our online experiment (N = 1,512) show that Senders transmit more information when negative externalities are larger and that Senders' decisions are largely driven by their own preferences towards the charity and their own use of information. We do not find evidence that Senders take the Receiver's preferences into account, as they largely ignore explicit requests for information, or ignorance, even if Receivers have the option to punish the Sender.

Keywords: Information avoidance; Lab experiment; Unethical behavior; Willful ignorance

JEL Codes: B41; C91; C93

1. Introduction

In many contexts, people have preferences over information and sometimes try to avoid it (Golman et al., 2017). Information avoidance often serves to protect cherished beliefs, for instance the protection of one's ego from bad feedback (Castagnetti & Schmacker, 2022) or the avoidance of bad financial news to reduce disappointment or stress (Sicherman et al., 2016). In particular, previous research has shown that some people try to escape responsibility for ethical decisions and maintain a good self-image by remaining uninformed about the consequences of their decisions (Dana et al., 2007; Grossman & van der Weele, 2017; Vu et al., 2023). Such willful ignorance may have important consequences for everyday consumption behavior, such as the decision to buy products that have adverse impacts on the environment or are manufactured in exploitative conditions (Ehrich & Irwin, 2005; Amasino et al., 2025).

Information avoidance also has an interpersonal side that has received much less attention. People often have information that is potentially inconvenient for *others*, and must decide whether to share

it. For instance, a vegetarian may ponder whether to give her carnivorous friends detailed information about the environmental costs associated with meat eating. In doing so, she may weigh several considerations. First, a concern for environmental consequences might motivate her to influence her friends' diets in the "right" direction. A second, more procedural reason to share would be to make sure her friends know the truth, whatever they end up doing. Finally, she may hold back information out of consideration for her friends' feelings. She may assess that transmitting information may make her friends feel judged, and even lead to confrontations that she may wish to avoid. Indeed, there is evidence that vegetarians and vegans sometimes experience backlash for sharing information about their diets, which causes some to keep a low profile (De Groeve & Rosenfeld, 2022; MacInnis & Hodson, 2017).

Other applications may occur within politics, organizations, or markets. For instance, politicians may have to decide whether to inform their voters about difficult trade-offs. Employees who have knowledge of organizational practices with negative external consequences, must decide whether to pass it up the decision-making chain. In buyer-seller interactions, sellers may voluntarily disclose ethical information about their products.

To study the trade-offs facing a sender of information, we designed an experiment that we call the "Button Game." The game involves two participants in the role of a Sender and a Receiver. The Receiver can press a large red button on the screen, which yields a bonus payoff of £1 for the Receiver, but may or may not degrade a fund destined for donation to a worthy cause. If the Receiver does not press the button, there are no additional payoffs for the Receiver or for the charity. The button is designed to be tempting; indeed, in the absence of specific information about the externality, virtually all Receivers in our experiment press the button and pocket 1.

Our primary interest is the decision of the Sender. Before the Receiver presses the button, the Sender can send information about the size of the externality at a small cost. In the *Baseline* treatment, Senders make multiple decisions for different sizes of the externality, where one of their decisions is randomly implemented. We find most Senders are willing to pay to send information, but only when externalities are relatively large. This indicates that some Senders trade off the payoff for the charity with the cost of sending. We also find evidence that personal preferences for inconvenient information, measured on a separate task, explain sharing. Finally, procedural concerns matter, as some senders share even if it does not change the Receiver's decision, and almost 30% of the Senders that share information say explicitly that this is the right thing to do.

To further investigate the Sender's willingness to accommodate the Receiver, we designed a treatment in which we vary the Sender's information about the Receiver's preferences. Before the Sender makes a decision, the Receiver can request either information or ignorance. We find no evidence that Senders respond to the Receiver's preferences, as neither the request for ignorance nor the request for information significantly affects information sharing. To reinforce the power of the request and mimic the possibility of conflict, we add a treatment with an option for the Receiver to punish the Sender by denying part of the Sender's participation payment. The threat of punishment does not make either type of request more effective, even though we observe some punishment by Receivers.

The key takeaway that emerges from our dataset is that sharing of inconvenient information is driven by the Sender's personal attitudes towards information and the externality. To the extent the findings from our stylized setting capture behavior outside the lab, the prevalence of sharing shows most people are motivated to share unethical information when it can have a significant impact. Nevertheless, the results also indicate the limits of sharing. The central role of Sender's own preferences for information in the decision to share suggests that sharing will be less prevalent for topics in which people are widely averse to information, like in the meat-eating example above (Epperson & Gerster, 2024; Onwezen & van der Weele, 2016).

Our paper contributes to a fast-growing experimental literature on information avoidance in ethical dilemmas (Dana et al., 2007; Grossman, 2014; Vu et al., 2023) and a smaller literature on how people share inconvenient information. Closest to our paper is Soraperra et al. (2023), who examine the demand and supply of willful ignorance in a market setup. Over multiple rounds, senders choose to release information or not, and decision-makers can choose to match with the sender they prefer. In this setting, senders suppress about 25% of inconvenient information on average, which correlates with their own preferences. However, the market setting is noisy, and there is not much control over the strategic incentives of the Senders or their beliefs about the decision-makers, preferences, making it hard to disentangle various explanations for information transmission and suppression. Another closely related study is Vellani et al. (2024). In an online experiment, they examine the motives of sharing potentially unpleasant information about monetary losses for the receiver. The results, which are in line with ours, show that participants use their own information-seeking preferences when deciding to share such information with others. Our study instead focuses on information sharing in the ethical domain.

A number of further studies look at information transmission. Lind et al. (2019) allow senders to force ethical information on decision-makers after they made their own decision to avoid or obtain information. They find that the option to be “overruled” by the sender results in more information seeking by decision-makers. Lane (2022) investigates a setting in which subjects can inform others about the externalities of their actions *after* they have taken a decision, so the information has no instrumental value but may reduce the happiness of the decision-maker. Most senders reveal information, despite the potentially negative impact on the receiver. In our paper, information has no instrumental value, but our finding that senders do not cater to the preferences of the receiver is in line with Lane (2022).

Our paper also has a link to research on paternalism. In particular, Ambuehl et al. (2021) find that people engage in an “ideals-projective” paternalism, where they assume their preferences are relevant to others and restrict others’ options accordingly. While senders in our study do not restrict any options, we do find that the sender’s own preferences for information and their evaluation of the charity are the main predictors of what they share with others.

The main contribution of our paper to this literature is to introduce a simple and portable setting to analyze the transmission of inconvenient information. We offer new evidence on the determinants of sharing decisions and the willingness of senders to accommodate the Receiver’s information avoidance.

2. Method and experimental design

The experiment consists of two tasks and a final survey. The first task measures participants’ preferences for information in an adaptation of the binary dictator game in Dana et al. (2007, DWK hereafter). The second and main task, a novel two-person game we call the “Button Game”, disentangles different motives to share information.

2.1. The DWK binary dictator game

Every participant played the binary dictator game, regardless of their role (Sender or Receiver) in the Button Game. The binary dictator game is inspired by the Hidden Information treatment proposed by Dana et al. (2007), with all participants acting as dictators and a charity as recipient, as in Lind et al. (2019). In this task, the participant has to choose between two options, namely, Option A and Option B, which have consequences for their payoff and for the donation to a charity, the Red Cross. The payoffs of Option A and Option B for the participant are £.6 and £.5, respectively. The payoffs for the Red Cross, instead, depend on the scenario: in the conflicting scenario A and B pay £.1 and £.5 to the charity; in the aligned scenario, the payoffs for the charity are flipped, with A and B paying



Fig. 1 Decision screen for the receiver in the Button Game. (a) Uninformed (b) Informed

£.5 and £.1, respectively. Participants are informed that each scenario is randomly selected with equal probability, and they can find out the realized scenario by clicking a *Reveal* button. Alternatively, they can select their preferred option directly, without knowing whether the payoffs for the charity follow the aligned or the conflicting scenario.

2.2. The Button Game

As the main task, we designed a two-person game in which a Receiver interacts with a Sender. The Sender possesses superior information about the consequences of the Receiver's action for a third party. The Sender can inform the Receiver before the latter chooses an action. We consider three variants of the game that manipulate how the two parties interact and define our treatments — namely, the *Baseline* treatment, the *Request* treatment, and the *Request + Punishment* treatment.

In all versions of the game, the Receiver has to decide whether to press a button; see Figure 1 for an example. The button is displayed for a total time of 30 seconds, during which the Receiver can press it. It was designed to be attractive to press: large and red. Pressing the button pays a bonus of £1 to the Receiver. In addition, it has consequences for a third party, the Red Cross, which range from +£.5 to -£2.5. Crucially, the Receiver has no information about the charity payoffs; neither about the actual value nor about the possible values. According to the instructions: “*Pressing the button also has consequences for the total amount donated to the Red Cross. These consequences can be either positive or negative, but you are not informed about them. They are concealed by ???*”.

Not pressing the button means that the Receiver will not get a bonus payment, but it ensures that the Red Cross will not be affected. To avoid Receivers pressing the button in order to speed up the experiment, subjects still have to spend the remaining time on the page. During this time, the button press cannot be reversed. On top of the button, we displayed the bonus of £1, alongside the payoff consequences of the charity – depending on the decision of the Sender.

The Sender is informed of the consequences for the charity, and this is common knowledge among the players. The Sender's task is to decide whether to share this piece of information with the Receiver before the latter makes their choice. The decision to pass information comes at a small cost of £.1 for the Sender.

In the experiment, we implemented the Sender's decision using the strategy method (see Appendix E for screenshots). Each Sender had to choose whether to share information for three negative impact levels (-£2.5, -£1.0, and -£.5) and one positive impact level (+.5). If the Sender decided to send information for a certain impact level and that level was randomly selected for implementation,




Fig. 2 Timeline of the different variants of the Button Game.

the Sender’s payoff for participation was reduced from £.5 to £.4. We tested understanding of these consequences with a comprehension question. A complete set of screenshots of all instructions and comprehension questions can be found in Appendix E.¹

The Button Game was designed to keep the strategic aspects of the Sender’s decision relatively simple. In particular, inspired by the sender-receiver game in Gneezy (2005), we kept the information to the Receiver about the payoff consequences for the charity down to a minimum. This feature encourages the Receiver to press the button in the absence of information. Moreover, since the Receiver does not know what kind of information the Sender can communicate, it limits the degree to which the Receiver can form beliefs about the externality in the absence of information, or form higher-order beliefs about the Sender’s intentions. This simplifies the analysis, where we will (mostly) abstract from such higher-order beliefs. It also simplifies the Sender’s decision problem, as she can assume that Receivers will press the button without information. To make sure that the Sender understands the decision environment of the Receiver, both Senders and Receivers start the game with an (unincentivized) practice round, where they can choose to press the button as an uninformed Receiver.

2.3. Timeline and treatments

Figure 2 shows the timeline of the Button Game and highlights the differences between the treatments. The software randomly allocates participants between the roles of Senders and Receivers. At the start of the game, all participants play a test round as an uninformed Receiver. In the *Baseline* treatment, the Sender moves first and decides whether or not to inform the Receiver. After the decision made by the Sender, the software randomly selects one of the four possible consequences for the charity. The information about these consequences is transmitted to the Receiver (or not), depending on the Sender’s decision. If it is transmitted, it is displayed at the top of the red button.

The Receiver thus decides whether to press the button with or without information about the consequences for the Red Cross, depending on the decision of the Sender for the selected consequence. The *Request* treatment extends the Button Game by adding a stage at the beginning where Receivers can either request information or ignorance about the payoffs for the charity. The Receiver selects from two pre-specified message options: There is no option not to send a request message. Finally, the *Request + Punishment* treatment extends the Request treatment by adding a stage at the end. In this final stage, the Receiver chooses whether to confirm or cancel the £4 bonus payment of the Sender. In the experiment, this decision was neutrally framed as “a final choice” to avoid normative connotations related to the word “punishment.” Finally, we administer a closing questionnaire.

¹The binary nature of the sending decision and the use of the strategy method may induce some experimenter demand effect by suggesting that sending is important, at least for some externality levels. We cannot test whether these elements affect senders’ decisions, but since they are kept constant in the experiment, they should not affect our comparisons between treatments.

2.4. Hypotheses

Here, we explain how we interpret the treatment differences and discuss our hypotheses. We preregistered the hypotheses prior to data collection.² Our hypotheses are not based on a formal model, but on a reasoned assessment of how subjects understand the game.³

Before diving into the hypotheses about the Sender's behavior, we briefly discuss what we expect for the Receiver. For the time being, we assume that these expectations reflect the Sender's beliefs about Receivers. As for the Receivers, we expect that virtually all of them press the button when uninformed, given that they earn a sure £1 bonus and the consequence to the Red Cross is ambiguous and possibly positive. Uninformed Receivers may also infer that the lack of information from the Sender meant that the externality was small or even positive, although the lack of precise information about the outcome means that they cannot form a precise Bayesian posterior.

When informed, instead, we expect that some of the Receivers will decide not to press the button to avoid generating harm to the charity. Moreover, we expect the likelihood of pressing the button to decrease weakly with the size of the consequences. Intuitively, if someone is willing to give up £1 to avoid a given level of consequences, the same person should also be willing to give up £1 to avoid more serious consequences.

Since our main interest is in the Sender's decision to share inconvenient information, we will focus only on Senders' choices for the negative consequences for the Red Cross, namely, we (mostly) ignore Sender's decisions for the £5.⁴ Moreover, we expect Senders to understand that in the absence of information, Receivers will press the button. This means that sending information about positive consequences is unlikely to make a difference to the outcome for the charity, although it may help the Receiver feel better about her choice.

For each sender, we define a "sender-index" that measures the point at which consequences for the Red Cross become too large not to share information. The index ranges from 0, when the Senders do not share information for any negative consequence, to 3, when the Sender shares information for all negative consequences. An index of 1 identifies those Senders that share information only for the most extreme (-£2.5) consequence, and an index of 2 identifies those Senders that share information for the -£1.0 and -£2.5 consequences but not for the least extreme (-£.5) consequence.

The *Baseline* treatment measures whether Senders have preferences for sharing inconvenient information of the Senders that are strong enough to overcome the small cost of sharing. As mentioned in the introduction, such preferences could depend on various motives, for example, (1) a concern for the charity, (2) procedural reasons like the belief that Receivers ought to make an informed choice, or (3) the desire to help the Receiver, combined with a belief the Receiver would like to be informed. Accordingly, our first hypothesis is that a non-negligible fraction of Senders decides to share inconvenient information.

Hypothesis 1. Senders send inconvenient information about the charity to their partners and are more likely to do so as the externality becomes more negative.⁵

To hypothesize the impact of the *Request* treatment, we consider both requests for information and for ignorance. The former request is straightforward to interpret: The main reason Receivers would like information is to decide whether to push the button. A request for information is, therefore, a signal that the information is likely to be used by the Receiver. Since sending helps both the Receiver

²For the preregistration, see Appendix F or https://aspredicted.org/X8Y_Q7T.

³In particular, we cannot analyze communication as a fully Bayesian game, as we did not tell the Receiver the possible outcomes for the charity nor the probabilities associated with these outcomes.

⁴The main reason for including a positive value is dictated by the need to truthfully tell the Receiver that consequences could be either positive or negative.

⁵The second part of this hypothesis was not preregistered but is implied by our use of the "sender-index," as explained above.

and the charity, we thus expect the Sender to increase the likelihood of sending information compared to the *Baseline*.

The effect of a request for ignorance is more complex. First, it may change the Sender's beliefs about the impact of information on the charity. The request may be a signal that the Receiver will not use the information, which may make the Sender less willing to send it. There is a caveat to this reasoning, however: The literature on moral wiggle room shows that a sizable fraction of subjects who choose to avoid information would nevertheless use it when they are confronted by it (Dana et al., 2007; Grossman & van der Weele, 2017). To the extent the Sender anticipates this, she may still perceive the potential impact on the charity to outweigh the cost of sending. To better understand how requests change beliefs, we therefore measure the Sender's beliefs about the Receiver's action in each condition. In addition, if the Sender cares about helping the Receiver, who expressed a wish for ignorance, one would expect the Sender to be more likely to suppress information.

Taken together, these considerations lead us to expect that Senders' decisions follow the direction of the request:

Hypothesis 2. Relative to the Baseline treatment, the likelihood of sharing inconvenient information increases with a request for information and decreases with a request for ignorance.

Finally, the comparison of the *Request + Punishment* is meant to amplify the strength of the *Request* treatment. In the *Request + Punishment* treatment, the Receivers can actually harm the Senders when they are unhappy about the provided or hidden information. Since Receivers can affect the Senders' payoffs, we expect Senders to follow the request of the Receiver more often. Furthermore, to the degree that the request affects the perceived cost of sharing information, the *Request + Punishment* treatment provides a measure of the cost sensitivity of the supply of inconvenient information.

Hypothesis 3. The possibility of Receiver punishment amplifies the impact of the requests on the likelihood of sharing information.

Along with hypotheses about the Sender's behavior, we derive secondary hypotheses regarding the impact on the overall welfare of the charity. Based on the previous hypotheses about Receivers' and Senders' behavior, we expect that Receivers requesting information are motivated by a willingness to avoid harming the charity. Therefore, the nature of the request, when obliged by Senders, will be correlated with the final outcome for the charity. Specifically, we expect the following:

Hypothesis 4. A request for information is associated with higher earnings for the charity and a request for ignorance with lower earnings for the charity. These effects are amplified in the punishment treatment.

2.5. Procedure

The experiment started with the binary dictator game, followed by the Button Game and the final survey. In the Button Game, participants were matched in pairs by the software, which meant that they had to wait for another player to join. If no other player appeared within 5 minutes, the software moved on to the end of the experiment, and the bonus payment was based on the results of the binary dictator game. When a match was possible, players were randomly assigned to the role of Sender and Receiver to start the Button Game.

After reading the instructions, both Senders and Receivers faced a practice round to experience the decision of the Receiver button page. In the practice round, no information about the consequences for the charity was communicated (see panel (a) of Figure 1). After the practice round,

Senders had to state their beliefs about the number of people pressing the button by moving a slider from 0 to 100 (“we ask you to think of 100 participants choosing as player A and give your best guess about how many of these chose to press the button”). To keep the game simple and payment quick, beliefs were not incentivized. Next, Senders answered a short set of comprehension questions. No comprehension questions were asked to the Receiver due to the simplicity of the button task.

At the end of the Button Game, Senders completed a belief elicitation page where they were asked to guess the likelihood that their Receiver pressed the button for each possible consequence, again unincentivized and on a scale from 0 to 100. In the *Request + Punishment* treatment, Senders were further asked to guess the likelihood of punishment. Receivers were also asked about their belief of other players A pressing the button. This page was identical to the Sender’s first belief elicitation page, but it was placed after the Receiver’s own choice to avoid spillover effects. At the very end of the experiment, all participants completed a demographics questionnaire, which included some open questions about their motivations in the Button Game and a 10-point slider to indicate how much they identified with the Red Cross (inspired by Ariely et al., 2009). Finally, each player was shown an overview of payoffs and was informed about the task that was randomly selected for the bonus payment.

The main study was run on Prolific in November 2022, where 1,796 participants started the study. Seventy-one of them dropped out before starting the DWK task, 5 could not be matched with another player, 171 dropped out during the DWK task, and 37 participants finished the study without a partner, leaving $N = 1,512$ responses (84.2% completion rate) for the analysis ($n_{\text{Baseline}} = 302$, $n_{\text{Request}} = 610$, $n_{\text{Request+Punishment}} = 600$). Due to practical constraints of the live matching into Sender-Receiver pairs, the treatments were run sequentially. To account for time effects, we started each treatment session approximately on the same time of day. All participants gave informed consent before participation. Participants were rewarded a £1.3 show-up fee plus the bonus earned in one of the two tasks, which was randomly selected at the end of the experiment. On the first page of the study, participants were informed about the payoffs to the Red Cross.⁶ We did not inform participants about the size of the original fund (which was £100). The experiment was programmed and data was collected via oTree software (Chen et al., 2016). The analysis code can be found at <https://osf.io/download/d3zcyj/>.

Several days after the end of the main study, all participants who completed at least the Dictator game (in one of the pilots or the main study) were messaged⁷ via Prolific with a proof of the donation to the charity.

3. Results

3.1. Preliminaries

Before analyzing Sender behavior, we first conduct a randomization check, and verify whether some key assumptions about the way both Receivers and Senders approach the game.

⁶The experimenters have prepared a fund to donate to the Red Cross at the end of the experiment. Your decisions may affect the size of this fund, and can either increase or decrease the total donations to the Red Cross. These donations are real, as our ethical approval does not allow us to deceive participants. A proof of the charity donation will be available upon completion of the experiment.

⁷Dear participant, In [month] 2022, you participated in our decision-making experiment on Prolific. As part of this experiment, we scheduled a donation to the British Red Cross. Based on the decisions in the experiment, positive and negative payoffs could be collected for the Red Cross. We would like to inform you that the donation to the Red Cross has been made. You can find the donation receipt and more details here: <https://figshare.com/s/d684b47812a2585174f4> Thanks again for your participation. You do not need to respond to this message. The researchers.

Table 1 Descriptive statistics by treatment

	Baseline	Request	Req. + Pun.	Overall	p-value
Gender (%)					.054
Female	136 (45)	316 (51.8)	323 (53.8)	775 (51.3)	
Male	161 (53.3)	291 (47.7)	273 (45.5)	725 (47.9)	
Non-binary/not say	5 (1.7)	3 (.5)	4 (.7)	12 (.8)	
Age in years (SD)	39.5 (12.9)	39.1 (13.)	39.0 (13)	39.1 (13)	.825
Monthly income (%)					.209
< £999	29 (9.6)	55 (9)	47 (7.8)	131 (8.7)	
£1,000-£1,999	68 (22.5)	156 (25.6)	130 (21.7)	354 (23.4)	
£2,000-£2,999	80 (26.5)	175 (28.7)	162 (27)	417 (27.6)	
£3,000-£3,999	61 (20.2)	103 (16.9)	123 (20.5)	287 (19)	
> £4,000	51 (16.9)	88 (14.4)	88 (14.7)	227 (15)	
Rather not say	13 (4.3)	33 (5.4)	50 (8.3)	96 (6.3)	
Identify with charity ^a (SD)	0.2 (2.9)	.2 (2.7)	.2 (2.8)	.2 (2.8)	.947
Browser type Desktop (%)	281 (93)	566 (92.8)	517 (86.2)	1364 (90.2)	<.001
Reveal in DWK (%)	119 (39.4)	256 (42)	223 (37.2)	598 (39.6)	.232
Press btn in the test round (%)	292 (96.7)	579 (94.9)	572 (95.3)	1443 (95.4)	.477
Observations	302	610	600	1512	

Notes: The table reports the means for the continuous and the counts for the categorical variables with, respectively, SD and percentages in parentheses. ^a Response to the question *How much do you identify with the charity Red Cross?* ranging from -5 = not at all to 5 = very much. The column “p-value” reports the results of a test comparing the different treatments. A chi-squared test is used for categorical variables and an Anova for the continuous variables.

3.1.1. Randomization check

Due to the live matching procedure of the Button Game, treatments were run sequentially on the Prolific platform. It may be the case that different user groups log into Prolific at different times and days of the week. Table 1 provides summary statistics about the participants’ demographics and other variables. The table allows us to assess the quality of the randomization across treatments. Overall, the sample is balanced regarding age, income, identification with the charity, button pressing in the practice round, and most importantly, own preferences for information (measured by the decision to reveal in the binary dictator game). Gender distribution (more women in the *Request + Punishment* treatment) and the device used (more mobile devices in the *Request + Punishment* treatment) are slightly unbalanced across treatments. To control for such differences, we added gender and device type as covariates in all further analyses.

3.1.2. Receiver’s behavior

We check whether Receiver behavior broadly aligns with our assumptions. Indeed, almost all Receivers press the button when uninformed (96.4%; $n = 364$) across all treatments. As mentioned in the hypothesis section, this high rate is unsurprising, given the monetary payoff of pressing the button and the absence of any information about the charity.

Moreover, all 61 Receivers that saw good news – namely, saw that the button increased the donation by an additional £.5 – pressed the button. Finally, the likelihood of pressing the button decreases with the severity of the negative consequence for the charity: 73.8% ($n = 107$) of the informed Receivers clicked the button when the consequences were -£.5, 66.1% ($n = 112$) when they were -£1, and 51.8% ($n = 112$) when they were -£2.5. This shows that, overall, the behavior of Receivers is

Table 2 Senders' beliefs about Receivers' button pressing, by consequence and treatment.

	<i>Dependent variable: Beliefs about Receivers' button pressing</i>			
	All tmts (1)	Baseline (2)	Request (3)	Req. + Pun. (4)
consequence+.5	5.344*** (.967)	4.715° (2.496)	4.593** (1.426)	6.423*** (1.511)
consequence-.5	-23.134*** (.981)	-20.808*** (2.132)	-24.636*** (1.517)	-22.777*** (1.608)
consequence-1.0	-31.354*** (1.075)	-26.596*** (2.360)	-33.272*** (1.631)	-31.800*** (1.776)
consequence-2.5	-40.458*** (1.210)	-36.172*** (2.671)	-42.466*** (1.868)	-40.573*** (1.968)
Observations	3,780	755	1,525	1,500
R ²	.465	.409	.490	.469
Adjusted R ²	.331	.258	.361	.334

Notes: Dependent variable: Response to the statement *I believe ... in 100 players will press the button*. Reference category: uninformed. Linear model with individual level fixed effects and heteroscedasticity robust standard errors in parentheses ($^{\circ}p < .10$; $^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$).

in line with our predictions, suggesting that they trade off the consequences for the charity with the cost of sharing.⁸ We will investigate the behavior of Receivers in more detail in [Section 3.6](#).

3.1.3. Sender's beliefs

Before we study the Sender's decision in detail, we first verify key Sender beliefs about the Receiver's behavior. In particular, to interpret the decision to share bad news as an attempt to help the charity, it must be true that Senders believe that sharing bad news leads indeed to a lower likelihood of pressing the button. Before making any decisions, Senders believe on average that 80.0% (SD = 18.1) of the Receivers press the button when not informed about the consequences. [Table 2](#) regresses Senders' beliefs that Receivers will press the button, conditionally on being informed, for the different possible consequences. It reveals that Senders' beliefs about Receivers pressing the button decline as the severity of negative consequences increases, compared to uninformed Receivers across all treatments. By contrast, Senders expect Receivers to press the button about 5 percentage points more often when they are informed about positive consequences. This shows that, on average, Senders (correctly) believed that sharing information would be effective, and more so when the externality was more negative. In [Section 3.5](#), we provide more details on the role of Sender beliefs in decision-making.

3.2. Information sharing in the baseline treatment

On the aggregate, Senders' decisions to share information increase with the size of the consequence. A positive consequence of £.5 is shared by 32.5% of Senders in the Baseline treatment. Sending information about a positive consequence may reassure Receivers that the charity benefits from their decision to press the button. However, there is a cost of sending information. Since most Senders (correctly) expect uninformed Receivers to press the button regardless, this can explain why most senders did not share this information. We discuss Sender's motives to send positive information further in [Appendix B1](#).

Negative consequences of -£.5, -£1, and -£2.5 are shared by 40.4%, 57.6%, and 71.5% of Senders, respectively. The differences between these proportions are statistically significant (pairwise

⁸The pattern is also confirmed also when looking at the Receiver's behavior separated by treatment.

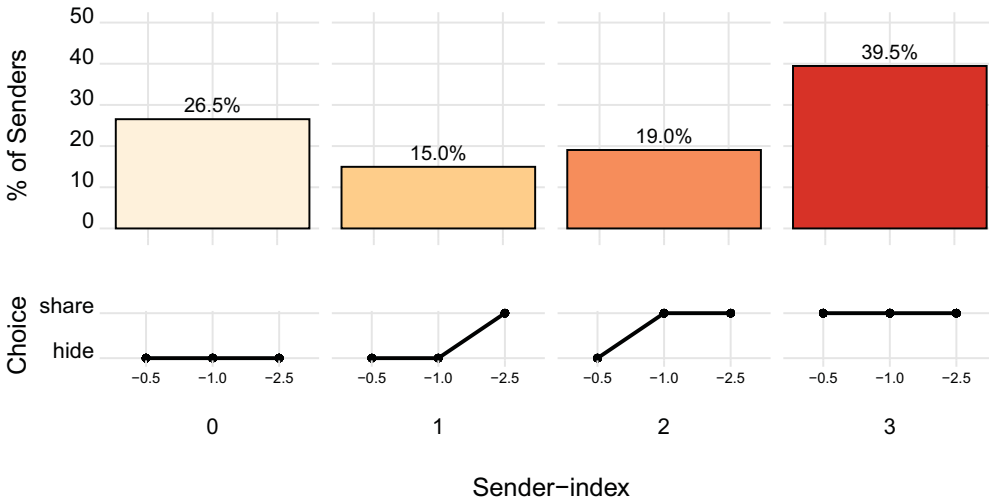


Fig. 3 Distribution of the sender-index in the *Baseline* treatment.

McNemar tests: consequence -0.5 vs. consequence -1.0 : $\chi^2(1) = 19.5, p < 0.001$; consequence -1.0 vs. consequence -2.5 : $\chi^2(1) = 17.4, p < .001$. Moreover, almost all Senders act consistently with a strategy where sharing small (negative) externalities implies sharing larger negative externalities. Only 40 out of 756 Senders decide to share information for less serious consequences and not for more serious ones.

This provides a rationale for our (preregistered) use of a “sender-index,” which reflects the smallest consequence for which the Senders decide to share information.⁹ Figure 3 shows the distribution of the sender-index in the *Baseline* treatment. This shows that there is substantial heterogeneity in the preferences for sharing information among our participants. On the one hand, 26.5% of the Senders never share information with the Receiver (sender-index = 0), and another 39.5% of Senders always share information about the consequences (sender-index = 3). The remaining Senders have intermediate preferences and share only when consequences are sufficiently negative (sender-index 1 and 2).

Overall, these numbers show that the majority of the Senders trade off the cost of sharing with the potential consequences for the charity, and provides support for Hypothesis 1.

3.3. The effect of requests

We now turn to the *Request* treatment, which allows us to investigate whether Senders take into account the preferences of the Receivers when sharing information. In this treatment, the majority of Senders (225; 73.8%) received a Request for information, while the rest (80; 26.2%) received a request for ignorance. Figure 4 (the three middle panels) shows the distribution of the sender-index across the various treatments. The Request and Punishment treatments are split by the nature of the request. According to our Hypothesis 2, we should observe an increase in the sender-index when the request is for information and a decrease when the request is for ignorance. However, the data do not show an increase in the frequency of Senders with higher sender-index values when transitioning from the left to the right panel in Figure 4. The average sender-index gives a similar picture, with an average of 1.73

⁹Following the preregistration protocol, we exclude the 40 non-monotonic participants from the analysis of the sender-index. Table B3 examines the pre-registered robustness check of a binary sender-index (including these 40 participants with non-monotonic sharing behavior). The appendix shows that results are robust to the inclusion of these participants. A detailed analysis of the sender-index can be found in Appendix A.




Fig. 4 Distribution of the sender-index by treatment.

when information is requested, of 1.65 when ignorance is requested, and of 1.71 when the request is not present. Indeed, a non-parametric Jonckheere-Terpstra trend test fails to reject the hypothesis of no difference in the sender-index across different requests ($z = .31, p = .377$).¹⁰

Regression evidence.

To examine the sharing decision more closely and with additional statistical power, we regress the sender-index on the treatments, as well as variables that measure the Senders' preferences for information and their identification with the charity. We also include control variables such as gender, age, income, type of device used, and the number of attempts to get the comprehension questions correctly. Since the sender-index is ordinal by nature, we employ an ordinal probit model to explore the correlation between such variables and the decision to share information.

Model (1) in Table 3 presents the results of the regression using the *Baseline* and *Request* data. It shows that requests for information have little effect, but requests for ignorance have a negative impact on the sender-index (compared to the Baseline without requests), although this is not statistically significant. Furthermore, sharing is positively related to how close the Senders feel to the charity, and whether they themselves revealed it in the DWK game. This result is highly statistically significant, and shows that preferences about the information one would like to have for oneself play an important role in sharing information with others.

To further understand the channels through which the request for ignorance affects the sender, we investigate whether senders are more likely to oblige when the request aligns with their own preferences for information. To test this, we run the same model restricting the data to those who remained ignorant in the DWK game (column 2) and those who informed themselves (column 3). The results show no significant interaction between the Sender's preferences and the request: While Senders who chose to remain ignorant are more likely to accommodate a request for ignorance, this effect is not statistically significant.

¹⁰Testing the more general assumption of a difference across distributions also does not support the idea that the request has an effect on the decision to share information. A χ^2 test cannot reject the null hypothesis of no differences in the distributions of the sender-index ($\chi^2(6) = 2.38, p = .882$)

Table 3 Ordered probit regressions of sender-index

	Dependent variable: Sender-index					
	(1)	(2)	(3)	(4)	(5)	(6)
Information preference (ref = Baseline)						
Request info	-.011 (.124)	-.111 (.162)	.085 (.202)	-.028 (.122)	-.072 (.157)	.023 (.197)
Request ignorance	-.130 (.172)	-.255 (.227)	.023 (.272)	-.150 (.169)	-.232 (.219)	-.063 (.273)
Request info under punishment threat				.096 (.130)	-.048 (.164)	.356 (.224)
Request ignorance under punishment threat				.083 (.167)	.097 (.228)	-.008 (.245)
Control variables						
Identify with charity	.990*** (.214)	1.234*** (.269)	.764* (.365)	.731*** (.167)	.884*** (.212)	.555 (.287)
Revealed in DWK	.257* (.119)			.323*** (.094)		
Log likelihood	-522	-312.5	-203.6	-842.8	-523.1	-311.9
Pseudo R ² (McFadden)	.038	.048	.028	.029	.028	.026
Covariates	Yes	Yes	Yes	Yes	Yes	Yes
Revealed in DWK		No	Yes		No	Yes
Observations	412	244	168	667	402	265

Notes: Ordinal probit model of the sender-index. Covariates suppressed for brevity: gender, age, income, browser type, comprehension questions. Models 1, 2, and 3 include all participants in the Baseline and Request treatments. Models 4, 5, and 6 include all participants across all treatments. Robust standard errors in parentheses (* $p < .10$; ** $p < .05$; *** $p < .01$; **** $p < .001$).

3.4. The effect of adding punishment

The *Request + Punishment* treatment allows us to test whether Senders stick to their preferences for sharing even when they risk punishment for not following the request (Hypothesis 3). Note that punishment rates were low, but not negligible. Requests were followed about half of the time, and deviations in responses to information requests were more likely to be punished (32%) compared to deviations after requests for ignorance (14.6%).¹¹ In line with the pre-registration, we test the null hypothesis that punishment does not change the pressure to follow the request of the Receiver against the alternative hypothesis that it increases the pressure to follow the request of the Receiver. Specifically, we test whether the threat of punishment increases the sender-index when information is requested and decreases the sender-index when ignorance is requested compared to the *Request* treatment.

As for the Senders’ decision in the *Request + Punishment* treatment, the left- and rightmost panels of Figure 4 show the distribution of the sender-index when a request for ignorance and for information are received, respectively. Visually, these distributions do not differ substantially from the ones observed in the *Request* treatment, which are reported in the second and fourth panel, respectively.

¹¹The distribution of requests observed in the punishment treatment is similar to the one observed in the treatment without punishment. In *Request + Punishment* treatment, 206 Senders (68.7%) received a request for information and 94 (31.3%) received a request for ignorance. 52 (17.3%) of the 300 Receivers punished the Sender for (not) responding to their request. Information requests were followed by 113 Senders and ignored by 93 Senders, of which 30 were punished (32.2%). Conversely, 48 Senders sent information when ignorance was requested, but only 7 of those were punished (14.6%). In a few cases, Receivers were punished when following the request for information (6 cases) or ignorance (9 cases).

Indeed, a one-sided Wilcoxon rank-sum test fails to reject the null hypothesis that punishment has no effect on the sender-index both when ignorance is requested ($p = .865$) and when information is requested ($p = .164$).¹²

We also use regressions to investigate whether requests combined with the threat of punishment induce different patterns from the baseline. Column (4) of Table 3 provides results that are in line with the graphical and non-parametric evidence: we do not observe any significant effect of requests when punishment is present. Columns (5) and (6) further reveal that the Sender's preferences for information do not interact with the request when the threat of punishment is introduced. All treatment coefficients in these regressions are statistically insignificant at the 5% threshold.

3.5. The role of sender's beliefs and motives

We previously showed that, prior to making a decision or learning about the requests, most Senders believed that Receivers will press the button when uninformed and that they are more likely to refrain from doing so when informed about the impact of externalities (see Table 2). In this section, we aim to better understand the Sender's motives in the decision to send information by analyzing text responses and stated beliefs. Appendix D provides the visual and statistical evidence corresponding to the claims in this section.

Motives in open-ended text responses.

In Appendix D1, we examine sender's motive to share elicited in open-ended text responses at the end of the experiment. Overall, these responses support the idea that a majority of Senders who revealed information did so out of a concern for the charity. Procedural motives are also prevalent, as almost 30% of Senders who send information indicates that it is the "right thing to do." A few Senders expressed concerns for Receiver's preferences, but this does not appear to be a dominant motive, in line with the muted impact of request. Overall, these results are in line with evidence in Lane (2022), which shows that most Senders will send information about negative externalities after Receivers have already made a decision, and the information is likely to have negative hedonic value. It is also consistent with Arrieta & Bolte (2023), who show that a majority of people think that having false beliefs is detrimental to a person's welfare.

Did the nature of the request affect Sender's beliefs in the Request treatment?

We find that the nature of the request does affect Sender's beliefs about the impact of sharing information on Receiver's behavior. Model (1) and (2) of Appendix Table D1 report regressions studying how the belief about the effect of sharing information – measured by the difference between the belief of pressing the button when informed and when uninformed – changes with the request and consequence levels. Focusing on the *Request* treatment and the largest externality level (-€ 2.5), Senders believe that sharing information when Receivers request it reduces the likelihood of Receivers pressing the button by almost 48 percentage points. This difference becomes smaller after receiving a request for ignorance: The estimated drop in the coefficient ranges from 9 to 11 percentage points, depending on the model, which is statistically significant at the 10% level.

While requests for ignorance decrease Sender's expectations that the Receiver will use the information, Senders still expect information sharing to reduce button pressing by more than 37 percentage points. Thus Senders expect that Receivers who prefer to remain ignorant may nevertheless refrain

¹²A Jonckheere-Terpstra trend test using all five combinations of treatments and requests does not reject the null hypothesis that the sender-index is not increasing in the pressure to follow the request ($z = .38, p = .352$). Similarly, a χ^2 test cannot reject the null hypothesis of no differences in the distributions of the sender-index across all five combinations of requests and treatments ($\chi^2(12) = 7.58, p = .817$).

from pushing the button when they are informed of the consequences. This expectation is partially correct (see Section 3.6) and helps explain the lack of response to requests.

Did the presence of punishment affect sender's beliefs?

In Models (5) and (6) of Table D1, we look at the difference in the Senders' beliefs about getting punished when sharing and when not sharing information. Regression results show that Senders do not perceive much difference in the likelihood of being punished when sharing or when withholding information, as the constants in the model are not significantly different from zero (although expected punishment is slightly higher after withholding information). Moreover, the nature of the request barely moves this expectation, suggesting that Senders were not sure how to interpret the request of the Receiver in this treatment. This may explain why punishment was not effective in enforcing responses to the request.¹³

Do beliefs explain sender's decisions?

To further understand the impact of beliefs on decisions, we regress the Senders' decisions on their beliefs that information makes a difference to the charity. We measure this as the decrease in the subjective Sender belief that the Receiver will press the button when informed, compared to being uninformed. We control for the treatments, the nature of the request, and the externality size.

The results of this exercise are presented in Appendix Table D2. There is a clear effect of Sender beliefs: For the case where the externality is -€ 2.5, namely, Model (1), an increase of 1 percentage point in the Sender's beliefs that information will sway the Receiver's behavior is associated with a statistically significant increase in the likelihood to send information of about .138 percentage point. Similar effects are observed in Models (2) and (3) for the other negative externalities.

Interestingly, the effect of the Sender's identification with the charity remains statistically significant, even when controlling for beliefs about the impact on the charity. One interpretation of this result is that Senders who care about the charity wish to signal to themselves that, regardless of whether Receivers act on the information, they have fulfilled their "duty" by providing information about the consequences. The effect of the Senders' own preferences for information, as measured by revelation in the DWK task, is also robust, underscoring the conclusion that Senders want others to have information that they value for themselves.

3.6. Receiver behavior and consequences for the charity

We now analyze the consequences of sharing information for the charity. Figure 5 shows the average Receiver impact on charitable donations in the different treatments. The left panel examines the aggregate effect of the treatment. The results show relatively small differences in the average payoff of the Red Cross. Statistically, we cannot reject the null hypothesis that the aggregated outcome is the same across the three treatments ($F(2, 753) = .87, p = .419$).¹⁴

In Hypothesis 4, we predicted that a request for information is associated with higher earnings for the charity and a request for ignorance with lower earnings for the charity. To visually evaluate the hypothesis, we split the results for the charity at the request of the Receiver. The right panel of Figure 5 suggests that Receivers who ask for ignorance cause more harm to the charity. A test that the distribution of charity outcomes is the same across all five groups shows a significant difference in the

¹³Moreover, the presence of punishment changes the Sender's beliefs about button pressing. As we discussed above, in the *Request* treatment, Senders correctly expected that information was less likely to have an impact on Receivers who requested ignorance. Models (3) and (4) in Table D1 and the left panel of Figure D2 show that this is no longer the case in the *Request + Punishment* treatment.

¹⁴Comparing the distribution of payoffs leads to the same conclusion ($\chi^2(8) = 4.55, p = .804$).




Fig. 5 Consequences for the charity. Average transfer of the Receiver to the charity fund (means and SE).

outcomes for the charity ($\chi^2(16) = 34.13, p = .005$). The OLS regressions in Table C1 show a negative effect of requesting ignorance (relative to Baseline) in both *Request* and *Request + Punishment* treatments, although the coefficients are only significant at $p < .1$. Moreover, we do not find evidence that punishment amplifies the effect of requests.

Receiver selection and response to information.

Above, we have shown that there are no statistical differences in Sender behavior in response to the request. Yet Figure 5 reveals that the impact on the charity varies with the request, suggesting that the Receivers' behavior correlates with the request they make. To understand this, we investigate the impact of both making a *request* for information and *actually* becoming informed on the behavior of Receivers. In Table 4, we regress the decision to press the button on dummies capturing information requested and received. To fully control for Senders' behavior, which varies with the size of the externality, we run separate regressions for each level of the externality. Moreover, we pool the data of the *Request* and *Request + Punishment* treatments to increase statistical power. Appendix Figure B1 shows the proportion of button clicks across treatment, type of request, and received message (again pooling the *Request* and *Request + Punishment* treatments).

The results show that uninformed Receivers have a similar behavior independently of their requests, as they universally press the button. Moreover, we observe that receiving information has a strong impact on pressing the button in the *Baseline* and that this impact increases with the size of the externality, as the estimated parameter of "Information received" gets smaller moving from Model (1) to Model (3). These results are in line with the raw data reported in Section 3.1.2.

When looking at the interaction of information preference and information received, we conclude that supplying information to those Receivers who request information has the same effect as supplying information in the *Baseline*. By contrast, supplying information to people who requested ignorance has a smaller impact than supplying information in the *Baseline*, principally because Receivers who request ignorance do not use the information as much. Nevertheless, Receivers who request ignorance are not wholly unresponsive to information, as they do not press the button as often as uninformed Receivers. This pattern is in line with previous research (Dana et al., 2007), and Senders seem to anticipate it (see Section 3.5).

Table 4 Receivers' button pressing by request and information obtained.

	Dependent variable: Button pressing		
	Cons. -2.5 (1)	Cons. -1.0 (2)	Cons. -.5 (3)
Constant	1.019*** (.061)	1.012*** (.056)	1.020*** (.057)
Information preference (ref = Baseline)			
Request info	.011 (.032)	.021 (.032)	.013 (.032)
Request Ignorance	.016 (.03)	.019 (.030)	.026 (.030)
Information (ref = Ignorance)			
Information received	-.546*** (.105)	-.393*** (.109)	-.192* (.100)
Interactions			
Request Info × Information received	.027 (.123)	.028 (.127)	-.126 (.117)
Request Ignorance × Information received	.362** (.134)	.202 (.130)	.101 (.114)
Observations	447	447	442
R ²	.374	.223	.181
Adjusted R ²	.358	.203	.160

Notes: Dependent variable: Button pressed. Covariates suppressed for brevity: revealed in DWK, age, income, browser type, comprehension questions. Linear model with heteroscedasticity robust standard errors in parentheses (* $p < .10$; ** $p < .05$; *** $p < .01$; **** $p < .001$).

In summary, our data show that a) information has a clear impact, as on average, Receivers want to avoid imposing negative externalities, and b) Receivers who request ignorance are less responsive to information and act more selfishly. Nevertheless, even among this last group, information does reduce button clicking.

4. Conclusion

We investigated the willingness to share inconvenient information in an online experiment. The key take-away that emerges from our dataset is that Senders are willing to pay to share inconvenient news out of concern for the otherwise negative consequences. The Senders' own preferences for information also play a significant role, thus showing that people share information that they consider in their own decision-making, in line with Vellani et al. (2024). We find little evidence that Senders try to cater to the preferences of Receivers. In particular, we do not find that they respond to explicit requests for ignorance or information, even when there is a threat of punishment. Indeed, Senders correctly anticipate that sharing information will still lead to better results for the charity, even if Receivers asked to remain ignorant.

If these results replicate in other settings, it implies that there is scope for inconvenient information to spread in society or organizations, as long as there are enough people who care about the affected third party. However, the results also point to the limits of information sharing. The fact that people share what they think is valuable for themselves suggests that people may end up in information silos. There is evidence that people dislike obtaining information that casts their behavior in a negative light (Golman et al., 2017; Vu et al., 2023). If social networks are characterized by homophily, namely, people interact with others who have similar preferences or behavior, this might lead to information

bubbles, in line with results in Soraperra et al. (2023). For instance, returning to the example in the introduction, there is evidence that meat eaters do not like to receive information about the negative consequences of meat (Epperson & Gerster, 2024; Onwezen & van der Weele, 2016). To the degree that carnivores seek each other out, they are unlikely to hear about the negative impacts of meat production.

There are a number of avenues for further research to address the limitations of the current study. One interesting extension would be to consider less anonymous interactions between senders and receivers, as in Foerster & van der Weele (2021). More generally, stronger forms of receiver punishment or opportunities for conflict may induce more self-censorship by senders. Second, one could look at different audiences: Perhaps people would be more motivated to share information with multiple receivers as the potential impact is bigger. One could also look at more extensive sharing networks to understand how inconvenient information spreads, perhaps testing predictions in Bénabou et al. (2020). Finally, one might look at various formats for information sharing, perhaps also including advice on the decision, which is the focus of Coffman & Gotthard-Real (2019).

Replication Packages. To obtain replication material for this article, <https://doi.org/10.17605/OSF.IO/MZPTY>.

Acknowledgements. We would like to thank Yves Le Yaouanq, Lenka Fiala, as well as conference participants at ESA Online Global 2021, TIBER 2021, ICSD 2022, and ESA World 2023 for their valuable comments. Sam Walsh, Britt van der Elsken, Ruben Bijl, and Paulina Schulte-Vels provided excellent research assistance. This research has received financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 865931) and the Dutch Research Council (NWO; Vi.Vidi.195.137 and 452.17.004).

Competing interests. The authors declare no conflict of interest.

References

- Amasino, D. R., Oosterwijk, S., Sullivan, N. J., & van der Weele, J. J. (2025). Seeking or ignoring ethical certifications in consumer choice. *Ecological Economics*, 229, 108467. <https://doi.org/10.1016/j.ecolecon.2024.108467>
- Ambuehl, S., Bernheim, B. D., & Ockenfels, A. (2021). What motivates paternalism? An experimental study. *American Economic Review*, 111(3), 787–830. <https://doi.org/10.1257/aer.20191039>
- Ariely, D., Bracha, A., & Meier, S. (2009). Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *American Economic Review*, 99(1), 544–555. <https://pubs.aeaweb.org/doi/10.1257/aer.99.1.544>
- Arrieta, G., & Bolte, L. (2023). What you don't know may hurt you: A revealed preferences approach. Available at SSRN.
- Bénabou, R., Falk, A., & Tirole, J. (2020). Narratives, imperatives, and moral persuasion. Working Paper 24 https://scholar.princeton.edu/sites/default/files/rbenabou/files/morals_september_15.pdf
- Castagnetti, A., & Schmacker, R. (2022). Protecting the ego: Motivated information selection and updating. *European Economic Review*, 142, 104007. <https://doi.org/10.1016/j.eurocorev.2021.104007>
- Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree: An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9(1), 88–97. <http://dx.doi.org/10.1016/j.jbef.2015.12.001>
- Coffman, L. C., & Gotthard-Real, A. (2019). Moral perceptions of advised actions. *Management Science*, 65(8), 3904–3927. <http://pubsonline.informs.org/doi/10.1287/mnsc.2018.3134>
- Dana, J., Weber, R. A., & Kuang, J. X. (2007). Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1), 67–80. <http://link.springer.com/10.1007/s00199-006-0153-z>
- De Groot, B., & Rosenfeld, D. L. (2022). Morally admirable or moralistically deplorable? A theoretical framework for understanding character judgments of vegan advocates. *Appetite*, 168, 105693. <https://doi.org/10.1016/j.appet.2021.105693>
- Ehrlich, K. R., & Irwin, J. R. (2005). Willful ignorance in the request for product attribute information. *Journal of Marketing Research*, 42(3), 266–277. <https://doi.org/10.1509/jmkr.2005.42.3>
- Epperson, R., & Gerster, A. (2024). Willful ignorance and moral behavior. Available at SSRN: 3938994.
- Foerster, M., & van der Weele, J. J. (2021). Casting doubt: Image concerns and the communication of social impact. *The Economic Journal*, 131(639), 2887–2919. <https://doi.org/10.1093/ej/ueab014>
- Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95(1), 384–394. <https://doi.org/10.1257/0002828053828662>
- Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1), 96–135. <https://doi.org/10.1257/jel.20151245>
- Grossman, Z. (2014). Strategic ignorance and the robustness of social preferences. *Management Science*, 60(11), 2659–2665. <http://pubsonline.informs.org/doi/abs/10.1287/mnsc.2014.1989>

- Grossman, Z., & van der Weele, J. J. (2017). Self-Image and willful ignorance in social decisions. *Journal of the European Economic Association*, 15(1), 173–217. <https://academic.oup.com/jeea/article-lookup/doi/10.1093/jeea/jvw001>
- Lane, T. (2022). Intrinsic preferences for unhappy news. *Journal of Economic Behavior & Organization*, 202, 119–130. <https://doi.org/10.1016/j.jebo.2022.08.006>
- Lind, J. T., Nyborg, K., & Pauls, A. (2019). Save the planet or close your eyes? Testing strategic ignorance in a charity context. *Ecological Economics*, 161, 9–19. <https://doi.org/10.1016/j.ecolecon.2019.02.010>
- MacInnis, C. C., & Hodson, G. (2017). It ain't easy eating greens: Evidence of bias toward vegetarians and vegans from both source and target. *Group Processes & Intergroup Relations*, 20(6), 721–744. <https://doi.org/10.1177/1368430215618253>
- Onwezen, M. C., & van der Weele, C. N. (2016). When indifference is ambivalence: Strategic ignorance about meat consumption. *Food Quality and Preference*, 52, 96–105. <https://doi.org/10.1016/j.foodqual.2016.04.001>
- Sicherman, N., Loewenstein, G., Seppi, D. J., & Utkus, S. P. (2016). Financial attention. *The Review of Financial Studies*, 29(29), 863–897. <https://doi.org/10.1093/rfs/hhv073>
- Soraperra, I., van der Weele, J., Villeval, M. C., & Shalvi, S. (2023). The social construction of ignorance: Experimental evidence. *Games and Economic Behavior*, 138, 197–213. <https://doi.org/10.1016/j.geb.2022.12.002>
- Vellani, V., Glickman, M., & Sharot, T. (2024). Three diverse motives for information sharing. *Communications Psychology*, 2(1), 107–120. <https://doi.org/10.1038/s44271-024-00144-y>
- Vu, L., Soraperra, I., Leib, M., der Weele, J. J. V., & Shalvi, S. (2023). Ignorance by choice: A meta-analytic review of the underlying motives of willful ignorance and its consequences. *Psychological Bulletin*, 149(9–10), 611–635. <https://doi.org/10.1037/bul0000398>

Cite this article: Mol, J. M., Soraperra, I., & van der Weele, J. J. (2025). Spoiling the party: Experimental evidence on the willingness to transmit inconvenient ethical information. *Experimental Economics*, 1–19. <https://doi.org/10.1017/eec.2025.6>