

The (In)Elasticity of Moral Ignorance*

Marta Serra-Garcia[†] and Nora Szech[‡]

August 2019

Abstract

Ignorance enables individuals to act immorally. This is well known in policy circles, where there is keen interest in lowering moral ignorance. In this paper, we demonstrate the relevance of demand elasticity to ignorance by showing that small monetary incentives can significantly reduce ignorance. We contrast monetary incentives with social norms, which have little impact on ignorance and actually increase ignorance in less moral individuals. Using a longitudinal design, we document that ignorance is persistent across moral contexts and through time. We propose and structurally estimate a simple behavioral model in line with our findings.

JEL Classification: D83; D91; C91.

Keywords: Information avoidance, morality, unethical behavior, social norms.

*We would like to thank Miguel Angel Ballester, Felix Bierbrauer, Jason Dana, Christine Exley, Eleonora Freddi, Russell Golman, Erin Krupka, Dorothea Kübler, Michael Kurschilgen, Jeff Naecker, Collin Raymond, Sally Sadoff, Nikolaus Schweizer, Charlie Sprenger, Josh Tasoff, Jeroen van de Ven, Joël van der Weele and the audiences at multiple seminars and conferences for helpful comments and suggestions. We would also like to thank Anke Greif-Winzrieth, Nicola Hüholt, and Giulia Maimone for excellent research assistance, and Sibille Hackel for administrative support.

[†]Rady School of Management, UC San Diego and CESifo. Email: mserragarcia@ucsd.edu.

[‡]Department of Economics, Karlsruhe Institute of Technology (KIT), Berlin Social Science Center (WZB), and CESifo. Email: nora.szech@kit.edu

1 Introduction

Ignorance allows individuals to engage more easily in questionable ethical behavior. For example, consumers do not want to know whether a favorite product benefits from child labor or causes damage to the environment, and they may even denigrate others who pursue this information (Zane et al., 2016). Likewise, many art collectors and museum managers ignore the origins of potentially stolen art work.¹ In business, it is not uncommon for managers to turn a blind eye to unethical behavior if those suspected of it are top performers (e.g., Rayner, 2012).² And in simple moral dilemmas, individuals avoid costless information about the consequences of their choices for others, thereby giving themselves “moral wiggle room” to act selfishly (e.g., Rabin, 1995; Dana et al., 2007; van der Weele, 2014; Bartling et al., 2014; Grossman, 2014; Grossman and van der Weele, 2017; Freddi, 2017; Golman et al., 2017, for a review).

An important question is whether moral ignorance can be easily reduced with monetary incentives or non-monetary incentives. This is a timely topic within policy circles, where there is hope that it may be possible to foster integrity through non-monetary tools (OECD, 2018), including “moral reminders” or “moral nudges.” But despite all of the attention and all of the interest, little is actually known about the effectiveness of different incentive types on moral ignorance.

In this paper, we study the elasticity of the demand for moral ignorance in response to both monetary and non-monetary incentives, using social norms as a sort of “moral nudge.” We show that monetary incentives, even if small, can significantly reduce moral ignorance. This elasticity is significantly stronger than in situations where no moral “good” is at stake. By contrast, social norms themselves have little impact on ignorance. In fact, people confronted with moral dilemmas often decide to avoid information, precisely because the information might pressure them to behave morally – and against their own self-perceived interests. While social norms increase the demand for information among those individuals who are likely to behave morally, they decrease it among those who are not.

¹For example, “The World Jewish Congress (WJC), a New York-based advocacy group, has criticized museums for waiting for artworks to be claimed by Holocaust victims instead of publicly announcing that they have suspect items” (Source: CNN.com, March 2, 2000).

²For example, Martin Winterkorn, former CEO of Volkswagen AG, argued that he would have stopped the emissions scandal if only he had known about it earlier. Yet, investigations suggest Winterkorn could have known already in 2007 (Source: USA Today, Jan. 19, 2017). In general, in large organizations, a manager’s ability to know about other individuals’ ethical behavior may be conveniently limited (Jackall, 1988; Dana, 2006).

The elasticity of moral ignorance is thus unlike that of information that is morally neutral and unlike that of moral behavior. It is also persistent across different morally-relevant contexts and through time, as we document using a novel longitudinal design. These findings therefore shed light on important new patterns regarding the malleability of ignorance, and have significant implications for policy interventions aimed at decreasing moral ignorance.

To study moral ignorance, we develop a new experimental paradigm, the “moral envelope game” (MEG). In this game, an individual chooses between an envelope that may – or may not – contain a donation to a worthy charitable cause, and a certain, private payment which the individual can pocket. We measure demand for information by allowing the individual to learn (or avoid learning) the envelope’s content before choosing between the envelope and the more selfish option. Opening the envelope or leaving it closed is monetarily incentivized, with different prices across a range of decisions.

To explain the demand for moral ignorance, we propose a parsimonious model in which rejecting the donation comes with a moral cost (DellaVigna et al., 2012; Andreoni et al. 2017). This moral cost may be absent (or much smaller) if the envelope is closed, that is, if the donation is uncertain (Exley, 2015). Indeed, in our study, individuals inclined to act selfishly preferred to remain ignorant about the envelope’s content. The model yields several predictions that we test experimentally, and can also be used to systematically examine the effects of different interventions.

This paper provides three main findings. First, we document significant elasticity to monetary incentives. Most subjects pay in order to leave the envelope closed and thus preserve their moral ignorance, as predicted in the model. This willingness to pay for moral ignorance can be calibrated using information demand in a morally neutral context, where the envelope contains a potential payment for the individual instead. In these contexts, with morality not a consideration, most individuals are willing to pay for information.

The demand curve for moral ignorance exhibits a sharp drop when moving from small monetary costs to small monetary incentives to acquire information. Such incentives can be an effective way of approaching people who have declined a request for a donation and who fear that more information will add to their guilt and their feeling of being pressured.

Second, we illuminate the complex interaction between social-norm information and desired behaviors, such as charitable donations. Existing research on social

norms has shown norms can often increase prosocial behavior (e.g., Cialdini and Goldstein, 2004; Schultz et al., 2007; Frey and Meier, 2004; Shang and Croson, 2009; Bicchieri and Xiao, 2009). Our paper investigates, for the first time, whether social-norm information is a powerful moral “nudge” to curb moral ignorance. In the MEG, avoiding information about the donation opportunity and acting selfishly is widely considered morally inappropriate, in contrast to obtaining information which is considered morally appropriate.

There is no question that social-norm information increases the likelihood of generous behavior among people who proactively seek out that information. The problem – in effect quantified by our study – is that many people do exactly the opposite. This raises important questions about the welfare effects of social-norm information (e.g., DellaVigna et al., 2012; Allcott and Kessler, 2019). Next to affecting the acquisition of information and donation decisions, social-norm information may increase the moral costs of declining a certain donation in selfish subjects and thereby increase the incentives of these subjects to avoid information, as we document descriptively and using a structural model. This result suggests that interventions based on social norms may come with little overall impact and the potential downside that moral costs may increase.

Third, we document that moral ignorance in the envelope game is predictive of ignorance in other decisions that have a moral component. Thus far, there has not been much to demonstrate the external validity of experiments measuring information avoidance in moral dilemmas. We address that gap, by looking at how our MEG subjects approached the controversial question of industrial livestock production. While consumers enjoy undeniable benefits from these production methods, they generally are not comfortable with the living conditions of the animals (te Velde et al., 2002; ASPCA 2016; BMEL 2016) and tend to avoid information on the topic (Onwezen and van der Weele, 2016).³ We decided to see how willing our MEG participants were to get information on these farming methods, and then to correlate their answers to their behaviors in the MEG game.

About a week after the MEG, the individuals who had participated were rewarded for correctly answering questions regarding industrial livestock production, and were offered the opportunity to watch an informative video beforehand. The results indicate moral ignorance is persistent across both tasks. To the best of our knowledge, this is the first experiment to provide suggestive evidence that preferring ignorance

³We follow Bandura (2016) in that avoiding unnecessary harm to humans, animals, and/or nature is of moral relevance.

in morally relevant decisions may be an individual trait that carries across moral contexts.

Our paper contributes to three strands of literature. First, a rich literature documents that individuals often seek excuses to avoid charitable giving and other prosocial behaviors. They avoid the charitable ask (e.g., DellaVigna et al, 2012; Andreoni et al., 2017; Exley and Petrie, 2018) or sharing decisions (Dana, Cain and Dawes, 2006; Broberg, Ellingsen and Johannesson, 2007; Lazear, Malmendier and Weber, 2012), and exploit risk (Exley, 2015) or potential poor charity performance (Exley, 2017) as excuses not to donate. Moreover, anecdotal evidence suggests individuals use news stories about high administrative costs and some charities’ high salaries – ignoring information about the charities’ performance – as excuses not to give at all.⁴ We study, for the first time, the elasticity of demand for information about a charitable opportunity.

Second, moral ignorance has been widely studied within the moral wiggle-room paradigm, in which an individual is unsure whether choosing the option yielding a higher monetary payment for herself hurts or helps another individual. Dana et al. (2007) were the first to show that individuals often avoid costless information about the consequences of their choices (costless meaning information that is free and a mere mouse click away) and are more likely to act selfishly as a result (see also Larson and Capra, 2009; Matthey and Regner, 2011; Feiler, 2014). Previous studies have considered the impact of a small (Grossman and van der Weele, 2017) or a larger cost (Cain and Dana, 2012) on the demand for ignorance.

We study ignorance in a paradigm that is very different from the one used to study moral wiggle room. While taking the own-payoff-maximizing option when one is ignorant is considered morally appropriate in the moral wiggle room paradigm (Krupka and Weber, 2013), it is not considered morally appropriate in the MEG. Rather, in our game, the morally appropriate choice is to either open the envelope, or choose the envelope, as it contains a 50% chance of a \$10 donation. Taking an outside payment of \$2.50, by contrast, is morally inappropriate.

In effect, we are providing a new setting to study information avoidance, in which ignorance is morally inappropriate and hence could be mitigated by non-monetary interventions such as social norms. We systematically elicit the demand curve for

⁴“Charities have brought skepticism on themselves in some cases by spending large percentages of donated funds on administrative costs and executive salaries. But this complaint is so commonly expressed now that it’s starting to sound like a dodge for not giving rather than a principled response to bad management at charities” (Source: The Globe and Mail.com, December 5, 2017).

moral ignorance, and compare the impact of monetary and non-monetary incentives on ignorance. Another difference is in the connection we make to an unrelated moral question, regarding animal welfare, which suggests that better understanding when ignorance is morally inappropriate can have wide-ranging implications in a variety of moral decision-making situations.

Third, recent surveys by Golman et al. (2017) and Hertwig and Engel (2016) demonstrate that ignorance does not only occur in morally relevant situations, but also in a variety of contexts, such as health (Oster et al., 2013; Ganguly and Tasoff, 2016), portfolio investment decisions (Karlsson et al., 2009; Sicherman et al., 2016), or work situations (Huck et al., 2018). Our findings contribute to this large and growing literature. We demonstrate that the moral relevance of the decision context is important to understand, and that moral ignorance exhibits unique features compared to other types of information demand.

The remainder of the paper is organized as follows. We next describe the setting in which information demand is studied and provide a parsimonious theoretical framework. Section 3 describes the design of the two main experiments and the additional experiment to measure the persistence of ignorance across contexts. Section 4 provides the results of the experiments, and section 5 discusses moral ignorance across different contexts. Section 6 concludes.

2 Avoiding Morally Relevant Information

2.1 The Moral Envelope Game (MEG)

We study moral ignorance in the following game. An individual is assigned an envelope that contains a \$10 donation to the Malaria Consortium, a non-profit organization fighting Malaria in sub-Saharan Africa, with $p = 0.5$. Otherwise (thus also with $p = 0.5$), the envelope is empty. The individual makes two decisions. First, she chooses whether to open the envelope or not. If she does not open the envelope, she chooses between receiving \$2.50 for herself and the uncertain envelope. If she opens the envelope, she first learns whether the envelope contains the \$10 donation or no donation. Then, she chooses between the envelope and receiving \$2.50 for herself.

For comparison, we analyze two morally neutral games. Self-10 is identical to the MEG, with the only difference being that instead of a donation, the envelope may contain \$10 for the individual. In Self-5, instead of \$10, the envelope may contain

\$5.

2.2 Theoretical Background

We solve the MEG and contrast predictions with those for the morally neutral games. We start with the case in which opening or leaving the envelope closed is costless, then turn to the case with direct monetary incentives for opening the envelope or keeping it closed. We also provide predictions on how social norms will influence the decision to open the envelope. We refer to choosing to take the private payment instead of the moral envelope as choosing the “selfish option,” and choosing the moral envelope as “donating.”

We assume utility takes the form of $u(x) = x^r$ with risk parameter $r > 0$. As we see, the predictions are independent of the risk parameter. Yet, they hinge on two parameters. First, it matters how much the individual values the donation of \$10, that is, which we capture by $\alpha \cdot 10$. The parameter α is thus a measure of altruism. Second, information demand depends on the moral cost associated with rejecting the certain donation of \$10. Research has shown that rejecting a certain donation opportunity can induce significant guilt or disutility, for example, from violating the social norm to donate (e.g., DellaVigna et al., 2012; Andreoni et al., 2017). We model this moral cost via a moral discount factor β . An individual who rejects the certain donation gets $\beta \cdot 2.5$ with $\beta < 1$, instead of 2.5. In other words, rejecting the certain donation feels morally worse than rejecting an uncertain one, and thus causes moral costs. Figure 1 provides the game tree.

When ignorance is costless, the main result is that moral discounting renders it dominant for individuals with low levels of altruism to leave the envelope closed and take the selfish option straight away. Only individuals with high levels of altruism open the envelope.

Proposition 1 *Consider the MEG when information is costless. In equilibrium, individuals with $\alpha < \frac{1}{4}$ keep the envelope closed. Individuals with $\alpha \geq \frac{1}{4}$ open the envelope. After opening the envelope, individuals donate unless the envelope is empty. After keeping the envelope closed, individuals choose the selfish option.*

When avoiding or obtaining information entails a cost, behavior depends both on the individual’s altruism and moral discounting. Suppose direct monetary incentives m_o for opening and m_c for keeping the envelope closed exist, as shown in the game




Figure 1: Game tree of the MEG when information is costless

tree in Figure 2. For simplicity, we focus on the case $r = 1$. In Appendix A, we also address cases of risk aversion and of risk lovingness.




Figure 2: Game tree of the MEG with direct monetary incentives

Proposition 2 *The equilibria of the MEG with monetary incentives $m_o, m_c > 0$ and $r = 1$ are as follows:*

- (i) *If $m_o - m_c \leq -\frac{5}{4}$, individuals keep the envelope closed regardless of α .*

(ii) If $m_o - m_c > 5 \left(\frac{1}{4} - \frac{\beta}{4} \right)$, individuals open the envelope regardless of α .

(iii) In the intermediate case,

$$-\frac{5}{4} < m_o - m_c \leq 5 \left(\frac{1}{4} - \frac{\beta}{4} \right),$$

a threshold value of $\alpha_t \in (\frac{\beta}{4}, \frac{1}{2})$ exists such that individuals open if $\alpha \geq \alpha_t$, whereas they keep it closed if $\alpha < \alpha_t$. The value of α_t is given explicitly through

$$\alpha_t = \frac{1}{4} - \frac{m_o}{5} + \frac{m_c}{5}.$$

In the experiments, $m_o - m_c$ ranges from -\$2 to \$2. Figure 3 depicts optimal information demand and donation behavior depending on the price of information and the level of altruism, for three cases of moral discounting. The level of altruism generally determines whether individuals open the envelope or leave it closed (see the transition from the green to the yellow area). If moral discounting is very strong, selfish individuals pay to remain ignorant. If it is mild, they may open the envelope and bear the moral costs when taking the selfish option (blue area). Individuals who leave the envelope closed most frequently take the \$2.50 payment, unless their altruism is very high ($\alpha > 0.5$), which we expected to be relatively rare.

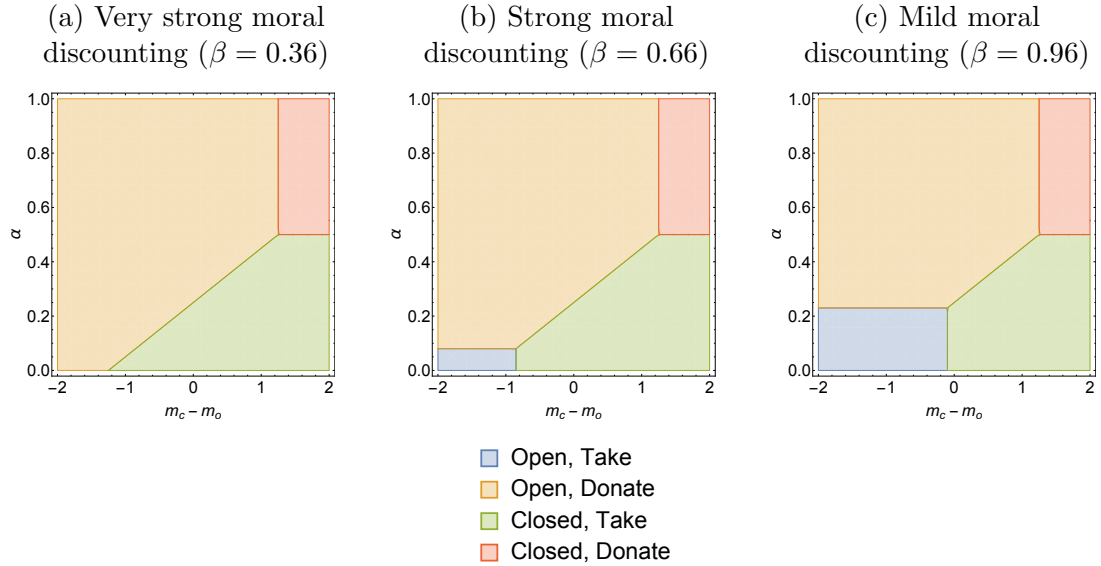


Figure 3: Optimal decisions for a risk-neutral individual with (a) very strong moral discounting ($\beta = 0.36$), (b) strong moral discounting ($\beta = 0.66$), and (c) mild moral discounting ($\beta = 0.96$), respectively.

From Figure 3, one can infer the demand curve for information depending on levels of altruism and moral discounting. To illustrate the demand curve directly, Figure 4 shows the demand curve for information under the assumption that α and β are continuously distributed, with α following a beta(1,3) and β a beta(3,1) distribution. We observe a non-differentiability around a price of information of zero. When ignorance is costly, β affects the slope of the demand curve. By contrast, when obtaining information is costly, α affects the slope of the demand curve.




Figure 4: Illustration of the demand for information, assuming α follows a beta(1,3) and β a beta(3,1) distribution.

Our analysis thus leads to the following hypotheses that will guide the data analyses. First, if individuals are rather selfish and exhibit moral discounting, they forego incentives to open the envelope. This way, they can take the selfish option without any moral discounting.

Hypothesis 1 *Subjects forego money in order to stay morally ignorant.*

To evaluate the impact of the moral context on ignorance, we compare ignorance in the MEG with the Self-10 and Self-5 games. In Self-10, the envelope contains \$10 for the individual. Altruism and moral discounting play no role, and hence the Self-10 game is equivalent to setting α and β equal to 1 in the MEG. The same is true for the Self-5 game, which is identical to Self-10 except that the envelope may contain \$5 instead of \$10. It corresponds to the case of $\alpha = 0.5$ and $\beta = 1$. In both cases, individuals pay for information. See Appendix A for details.

Observation 1 *In the morally neutral Self-5 and Self-10 games, subjects pay for information.*

In the MEG, the slope of the demand curve for information to the left of zero, when avoidance is costly, is determined by β , whereas the slope to the right, when information is costly, hinges on α . Therefore, we expect a non-differentiability of the demand curve at zero.

Observation 2 *In the MEG, the slope of the demand curve when avoidance is costly is different than the slope when information is costly.*

Next, we can examine heterogeneity in information demand and, as our analysis shows, whether individuals demand or avoid information depends on their level of altruism. Whereas individuals with low levels of altruism in the MEG pay to stay ignorant, those who are significantly altruistic demand information even if it is costly.

Hypothesis 2 *In the MEG, low levels of altruism correlate with a high willingness to pay for moral ignorance.*

As highlighted above, the MEG and the Self games are directly connected. Specifically, subjects in the MEG who display an α of 0.5 should behave as subjects in Self-5, whereas subjects in the MEG with an α of 1 should behave as subjects in Self-10.

Observation 3 *Subjects with very high levels of altruism in the MEG should behave as subjects in the Self treatments.*

In addition to the impact of monetary incentives, we study empirically the causal effects of providing information about social norms. We expected these norms to be in favor of opening the envelope and against taking the selfish option while leaving the envelope closed. In terms of our model, we expected that social norms would increase willingness to donate (WTD), leading to a higher level of altruism (α). Such an increase could lead to more demand for information. Yet, social norms could have a partially opposing effect if norms also increase moral discounting. If so, individuals with still low levels of altruism would pay even more to stay morally ignorant.

Hypothesis 3 *Social norms may increase demand for information in rather altruistic subjects, but reduce information demand in selfish subjects due to stronger moral discounting. If norm provision increases levels of altruism, an overall increase in information demand becomes likely.*

3 Experimental Design

In the following, we present the designs of two experiments examining the impact of monetary incentives (Experiment 1) and social norms (Experiment 2) on ignorance. In total, we analyze the decisions of 1,304 participants on Amazon Mechanical Turk (MTurk).⁵ The design of the two experiments is summarized in Table 1. The experiments were pre-registered. Throughout, we refer to analyses that were not pre-registered as explorations of the data.

Table 1: Experimental Design

Experiment	Treatments	Description
1	MEG	Donation uncertainty: \$10 donation with $p = 0.5$, \$0 otherwise
	Self-5	Payment uncertainty: \$5 payment with $p = 0.5$, \$0 otherwise
	Self-10	Payment uncertainty: \$10 payment with $p = 0.5$, \$0 otherwise
2	Norms	Elicitation of social norms regarding information demand
	NoNorm	MEG without norm information
	Norm-Avoid	MEG, avoidance (keep closed and take) is morally inappropriate
	Norm-Seek	MEG, seeking information (open) is morally appropriate

3.1 Experiment 1

In this experiment we study the effect of incentives on preferences for information. Our main treatment is the MEG, in which the envelope contained a \$10 donation with a 50% chance. Individuals had to choose between one of three options: take a \$2.50 payment, take the envelope, or open the envelope first. Hence, the decision setting contained no default choice (for an analysis of default effects, see Grossman, 2014). We varied the payment for opening the envelope from -\$2 to \$2. Specifically, each individual made nine independent decisions, with the following range of payments for opening the envelope: \$2, \$1, \$0.50, \$0.10, \$0, -\$0.10, -\$0.50, -\$1 and -\$2.⁶

To benchmark preferences towards morally relevant information against such preferences in a morally neutral context, we ran two control treatments (Self treatments).

⁵We used TurkPrime to run the studies and to re-invite participants to the follow-up task (Litman, Robinson and Abberbock, 2016). Further details are provided in Appendix C.

⁶Individuals knew that whether the envelope contained the donation or not independently varied across all nine decisions. To simplify elicitation, decisions were made one at a time, on separate screens, and the order of the questions always followed the same descending pattern of payments for opening the envelope. The instructions are presented in Appendix B.

In these, we replaced the uncertain donation with an uncertain payment for subjects themselves. We designed the Self treatments to match average willingness to pay for ignorance in the MEG with an equivalent willingness to pay for information in the Self treatments.⁷ The question we answer is, what dollar amount placed in the envelope in the Self treatments would result in a willingness to pay for information of the same magnitude as the average willingness to pay for ignorance in the MEG treatment? We first ran the Self-5 treatment in parallel with the MEG treatment. We learned that the willingness to pay for information in Self-5 was lower than the willingness to pay for moral ignorance in the MEG treatment. Thus, we ran the Self-10 treatment.

3.2 Experiment 2

In Experiment 2, we study social norms regarding ignorance. A large number of studies have shown that social norms affect individuals' behavior in an array of contexts (e.g., Cialdini and Goldstein, 2004; Schultz et al., 2007; d'Adda et al., 2018). In the context of donation behavior, injunctive-norm information, which describes how individuals should behave, can increase prosocial behavior (e.g., Bicchieri and Xiao, 2009; Hallsworth et al., 2017).

Thus far, no study has examined how information about social norms affects information demand. Philosophers have proposed the "Ignorance Thesis," which states that if an individual chooses to remain ignorant in a moral decision, she is culpable for acts that derive from it (Zimmerman, 1997; Rosen, 2003; Guerrero, 2007). If individuals broadly agree with this view within the MEG, we should find injunctive norms in favor of information demand.

We expected that in our setting remaining ignorant and choosing the selfish payment would be considered morally inappropriate, because it implies that the individual keeps \$2.5 while forgoing a 50% chance of a \$10 donation. At the same time, we expected that it would be morally appropriate to remain ignorant and donate the envelope, or to open the envelope first. We thus expected our setting to be substantially different from the moral wiggle room game in Dana et al. (2007) from a moral perspective. In that game, choosing the own-payoff-maximizing option increases the payoff of the decision maker by \$1. This choice can yield \$4 additional dollars to another subject with a 50% chance, or decrease it by the same amount with a 50% chance. Krupka and Weber (2013) find that ignorance while taking

⁷We thank Charlie Sprenger for suggesting this setup of the calibration exercise.

the own-payoff maximizing option is considered morally appropriate, more so than remaining ignorant and taking the option yielding a lower payoff.

We first ran the Norms treatment to elicit the perceived morality of information and donation decisions, using the method proposed by Krupka and Weber (2013). Subjects rated decisions as “very morally appropriate,” “somewhat morally appropriate,” “somewhat morally inappropriate,” or “very morally inappropriate.” Ratings were elicited for each of the nine information decisions, that is, for each cost of information (avoidance). They earned \$5 if their rating in a randomly drawn decision coincided with the most frequently chosen answer of the other participants in that treatment, and zero otherwise.

On average, more than 70% of individuals considered it inappropriate (somewhat or very inappropriate) to take the \$2.50 payment without demanding information first and more than 70% considered it (somewhat or very) appropriate to demand information. Also, more than 70% of individuals considered it morally appropriate to donate without demanding information. As we show, these beliefs did not vary strongly by the price of information. Hence, this invariability allowed us to provide to a new set of subjects simple information based on our norms elicitation.

We randomly allocated a new set of participants to one of three treatments. The first treatment was a NoNorm treatment, which was the same as the MEG treatment in Experiment 1. The second was a Norm-Avoid treatment, in which individuals were informed that over 70% of individuals considered taking the \$2.50 payment to be morally inappropriate, without seeking information first. The third treatment was a Norm-Seek treatment, in which individuals were told that over 70% of individuals considered seeking information to be morally appropriate.⁸ We hence tested two simple messages regarding social norms, one positively and one negatively framed, that could decrease moral ignorance. Each message was shown once, just before individuals started making their information decisions.

3.3 Follow-Up: External Validity of Information Preferences

An important question when measuring ignorance in an experimental task is whether such behavior is externally valid in a different incentivized setting. We are aware of no evidence regarding the predictive validity of experimental measures of ignorance

⁸The message shown to participants was “over 70% of MTurkers who evaluated the actions in this part of the study consider it morally inappropriate (appropriate) to choose the option “Get \$2.50” without revealing what the envelope contains first (“Reveal what the envelope contains” first).”

in a different context. To test the predictive power of the elicited informational preferences, we later measured ignorance in a different, morally relevant task.

We invited participants of our experiments to an unrelated work task between 7 and 10 days after they had participated in the experiments described above. The task consisted of answering questions about the living conditions of cows and their calves in conventional dairy production. We chose this topic because the willingness to improve living conditions of farm animals correlates with a higher moral and pro-social inclination (Albrecht et al., 2017). Even though many consumers buy products from intense animal farming, many state that they do not agree with the living conditions of animals involved (te Velde et al., 2002).

Participants earned a \$0.15 bonus if they correctly answered two questions about the treatment of cows and their calves in conventional farming. Before proceeding to the questions, they were offered the option to watch a 1-minute informational video. We study whether individuals who choose to remain ignorant in the MEG also choose to avoid watching the video.⁹

3.4 Experimental Procedures

3.4.1 Other Determinants of Ignorance

In all experiments, subjects first played the respective game in their treatment. Thereafter, we elicited several control measures of individuals' preferences.

As we saw in the Theoretical Background section, a subject's valuation of the donation (α) is particularly relevant for her decision on whether to remain ignorant. To measure this valuation, subjects participated in a task that elicited the monetary equivalent of a \$10 donation. Individuals made a series of eight binary choices in each task, choosing between a payment that increased from \$0.10 to \$10 and a \$10 donation. Their monetary equivalent, or willingness to donate (WTD), is measured as the maximum payment that individual was willing to give up instead of the certain donation.¹⁰

At the end, in a questionnaire, we elicited preferences for information (avoidance) according to the Monitors-Blunters Scale (Miller, 1987), moral preferences accord-

⁹In contrast to the MEG, here, repugnance (e.g., Roth, 2007) may play an additional role for avoiding information. Nevertheless, we observe a marked correlation across the two contexts.

¹⁰They also made choices when the donation was uncertain, deciding between a payment that increased from \$0.01 and \$5 and a 50% chance of a \$10 donation. Subjects knew that the computer either drew one of the two donation valuation tasks or the main part of the experiment for payment, and that one decision in this task would be implemented.

ing the Machiavellianism scale (Christie and Geis, 1970), as well as gender, age, education, and frequency of work in MTurk. The Monitors-Blunters Scale (Miller, 1987) is a well-established scale in psychology that measures information-seeking behaviors under threat. A higher score on the scale implies more information-seeking behavior. We test whether subjects who express more desire for information in the Monitors-Blunters Scale are less likely to remain ignorant in the MEG treatment. The Machiavellianism Score measures whether an individual considers herself attached or detached from moral and social norms (Christie and Geis, 1970). We test whether individuals classified as more Machiavellian have a lower willingness to remain ignorant in the MEG treatment.

In the Self treatments, we also elicited a control measure of subjects’ risk preferences. After subjects had completed the respective envelope game, we elicited their risk preferences and used the same questionnaire described above. Because these treatments included no mention of a donation opportunity, we did not measure subjects’ valuation of the \$10 donation to fight malaria. We provide detailed information on these measures and descriptive statistics in Appendix C.

3.4.2 Sample

Experiment 1 included 593 subjects, excluding inconsistent subjects as pre-registered. Of these, 294 participated in the MEG treatment, 147 in Self-5, and 152 in Self-10. In Experiment 2, we first conducted the Norms treatment, with 102 participants. We thereafter conducted the treatments NoNorm, Norm-Avoid and Norm-Seek at the same time and randomly assigned subjects to one of these three treatments. The treatments contain 200, 201, and 208 subjects, respectively. These sample sizes were chosen to be able to detect a \$0.15 change in willingness to pay for information in the presence of social-norm information (with an 80% power).

Because Experiment 2 was conducted two months after Experiment 1, we again elicited the behavior of individuals in the MEG treatment, labeled NoNorm treatment here, to control for any differences in the sample.¹¹ We find that, in contrast to Experiment 1, the share of female participants in Experiment 2 was significantly higher, 53.0%, compared to 45.2% in Experiment 1 (t -test, p -value < 0.01).¹² In line

¹¹We conducted a first smaller version of Experiment 2 that suggested results would be incomparable to the former MEG treatment run earlier. We hence conducted a larger study thereafter, and focus on these data. Including the smaller study does not change the conclusions.

¹²Age, educational attainment, and MTurk experience of participants did not differ (t -tests, p -value > 0.05).

with previous literature on gender effects in altruistic behavior (e.g., Andreoni and Vesterlund, 2001), we observe a higher monetary equivalent of the \$10 donation in this experiment, \$2.90, compared to \$1.90 in Experiment 1 (t -test, p -value < 0.01). We also observe a weaker preference to remain ignorant in the NoNorm treatment than in the MEG treatment in Experiment 1 (t -test, p -value < 0.01), consistent with our prediction (and finding) that a higher level of altruism is associated with a lower willingness to pay for ignorance. The analysis of Experiment 2 hence focuses on the treatment effect of providing information about social norms within this experiment.

On average, 86.3% of participants in the experiments completed the follow-up task in which we measured information demand about cows' living conditions to examine the persistence of ignorance across tasks.¹³ Because the relationship between ignorance in the envelope game and the follow-up task is qualitatively similar in the MEG treatment in Experiment 1 and in Experiment 2, we focus our analysis on Experiment 2, where we can also examine whether social norms had a long-run effect on ignorance.

Finally, a concern when running an experiment on ignorance, especially online, is that subjects remain ignorant in order to save time. Obtaining information involves only two additional clicks by the participant (selecting the envelope or the payment for herself, and moving onto the next question), which takes very little time, whereas direct incentives to obtain information can be very high (of up to \$2). Subjects who pay to remain ignorant do not finish earlier than those who do not.¹⁴ We also conducted a pilot experiment in a laboratory setting at UC San Diego. In this pilot we find comparable rates of information demand.

4 Results

4.1 Experiment 1: Price Elasticity of Ignorance

Figure 5 displays the share of individuals who demand information, that is, open the envelope, for each price of information, by treatment. The bottom black line depicts

¹³The return rate is 83.7% in the MEG treatment in Experiment 1, 91% in the NoNorm treatment in Experiment 2, and 87.6% and 84% in the Norm-Avoid and Norm-Seek treatments, respectively. Within Experiment 2, the difference in return rates between the Norm-Seek and NoNorm treatments is significant (p -value = 0.036).

¹⁴The relationship between willingness to pay for ignorance and decision time in the experiment is actually negative. We find that subjects who pay to remain ignorant take on average somewhat more time than those who pay for information (Spearman rank correlation coefficient $\rho = -0.1264$, $p = 0.03$).

information demand in the MEG treatment. We observe that 31% of individuals demand information when the cost of information is zero. This finding implies that the majority, that is, 69% of individuals, prefer ignorance when it is costless. Among these information avoiders, 65% are willing to pay at least 10 cents, and 57% are willing to pay more than 50 cents for moral ignorance. Further, 44% of them are willing to pay at least \$1, and 22% are willing to pay \$2 for ignorance. This result illustrates that, on one hand, introducing a small monetary cost of ignorance increases information demand significantly, from 31% to 55%, i.e., 24 percentage points. On the other hand, for a substantial proportion of subjects, preferences for moral ignorance are strong. On average, across all subjects in the MEG treatment, individuals pay 40 cents in order to stay morally ignorant. This amount is significantly different from \$0 (t -test, $p < 0.01$), and suggests moral costs of rejecting the certain donation exist, in line with Hypothesis 1.




Figure 5: Information Demand, by treatment

In the Self treatments, by contrast, we observe that individuals are, on average, willing to pay for information acquisition. The average willingness to pay for information in the Self-10 treatment is \$0.83, and \$0.29 in the Self-5 treatment. In both cases, this amount is significantly different from zero (t -test, $p < 0.01$). The will-

ingness to obtain information in the Self-10 is higher than in the Self-5 treatment ($p < 0.01$), in line with standard comparative statics of information preferences. Hence, in a morally neutral context, individuals pay to obtain information, in line with Observation 1.

The Self-5 and Self-10 treatments serve two further purposes. First, they allow us to calibrate the magnitude of moral ignorance. Assuming linearity, we find willingness to pay for ignorance in the MEG treatment is equivalent to willingness to pay to obtain information when the envelope contains approximately \$6 for the individual with 50% chance, that is, a Self-6 treatment. A figure illustrating this equivalence is provided in Appendix D. Second, the Self treatments allow us to identify the causal effect of the moral context on ignorance. Our results indicate roughly half of ignorance at a cost of \$0 is specific to the moral relevance of the situation (68.7% in the MEG, compared to 34% and 25% in the Self-5 and Self-10 treatments, respectively).

To examine the elasticity of ignorance with respect to monetary incentives, we estimate the slope of the demand curve. According to Observation 2, we can expect differences in the slope of the demand curve at zero, when information becomes costly to obtain instead of being costly to avoid. We conduct an exploratory analysis of the differences in price sensitivity in the MEG, depending on whether information is costly to avoid or to obtain, and compare it with that in the Self treatments in Table 2. This table presents the results of linear probability models on the decision to demand information, as a function of the price of information. The regression includes an indicator variable for costly information, that is, when prices are strictly positive, to examine how the demand curve changes around a price of information of 0. To allow for the slope of the demand curve to vary when information is costly relative to when it is costless, the regression also includes an interaction term between the indicator for costly information and the price of information.

The results shown in columns (1)-(3) of Table 2 display the relationship between price and information demand in the MEG, Self-5 and Self-10 treatments, respectively. Column (4) shows the effect of price is different in the MEG treatment, compared to the Self treatments, in three ways. First, in the MEG treatment, we observe a pronounced drop of 27.51 percentage points when avoidance is no longer costly. This drop is significantly different from the Self treatments which show more demand for information and a smaller drop of 6.7 and 6.1 percentage points. The drop in the MEG treatment is consistent with suggestive evidence in Grossman and van der Weele (2017), who find information demand drops to zero when a 0.10 Euro

Table 2: Demand for Information Across Domains

<i>Treatment:</i>	(1) MEG	(2) Self-5	(3) Self-10	(4) All
Price (of Information)	-0.2159*** (0.0130)	-0.1124*** (0.0170)	-0.0902*** (0.0154)	-0.2159*** (0.0130)
Costly Information	-0.2751*** (0.0234)	-0.0670** (0.0271)	-0.0613** (0.0242)	-0.2751*** (0.0234)
Costly Information X Price	0.1581*** (0.0185)	-0.1499*** (0.0347)	-0.1036*** (0.0305)	0.1581*** (0.0185)
Self-5				0.2371*** (0.0469)
Self-10				0.3398*** (0.0417)
Price X Self-5				0.1035*** (0.0213)
Price X Self-10				0.1257*** (0.0202)
Costly Information X Self-5				0.2081*** (0.0358)
Costly Information X Self-10				0.2138*** (0.0337)
Costly Information X Price X Self-5				-0.3081*** (0.0392)
Costly Information X Price X Self-10				-0.2617*** (0.0356)
Constant	0.4452*** (0.0266)	0.6823*** (0.0387)	0.7851*** (0.0322)	0.4452*** (0.0266)
Observations	2,646	1,323	1,368	5,337
R-squared	0.3117	0.2301	0.1641	0.3166
Nr. of subjects	294	147	152	593

Notes: This table examines the impact of price on information demand in the MEG, Self-10 and Self-5 treatments, using linear probability models. The dependent variable takes value 1 if the individual demands information (opens envelope). Robust clustered standard errors shown in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

incentive is offered to avoid information, in an explorative session with 10 subjects in the moral-wiggle-room paradigm of Dana et al. (2007). Second, when individuals are paid to acquire information, price sensitivity is stronger in the MEG treatment than in the Self treatments. This difference is explained by the fact that most subjects always demand information in the Self treatments, whereas many subjects exhibit a preference to remain ignorant in the MEG treatment. Third, price sensitivity is weaker in the MEG treatment than in the Self treatments when individuals have to pay for information. Because only a few subjects have a preference to pay

at all for information in the MEG treatment, seeing a smaller elasticity here may not be surprising.

4.1.1 Structural Estimation of Preferences for Ignorance

Because we find evidence of some ignorance in all treatments, one may wonder whether some subjects made mistakes. Therefore, we conduct exploratory analyses that allow for noise in subjects' behavior, and estimate the extent of altruism and moral discounting structurally.

Using a nested logit model (see, e.g., Cameron and Trivedi, 2005, Ch. 15) for the MEG treatment, we estimate average altruism (α) and moral costs (β), as well as the CRRA parameter (r), as detailed in Appendix E. The results are shown in column (1) of Table 3. We find the estimated α is 0.21 (*s.d.* = 0.01) and the estimated β is 0.90 (*s.d.* = 0.02), which is significantly smaller than 1 ($p < 0.01$), consistent with significant moral discounting.

Table 3: Estimation of Preferences for Ignorance

	(1) MEG	(2) Self-5	(3) Self-10
Altruism parameter $\hat{\alpha}$	0.2101 (0.0112)		
Moral cost parameter $\hat{\beta}$	0.8982 (0.0256)		
Risk aversion parameter (CRRA) \hat{r}	1.0289 (0.0258)	0.9716 (0.0335)	0.7771 (0.0291)
Loglikelihood	-3297.09	-1109.09	-1153.51
Observations	2,646	1,323	1,368
Nr of subjects	294	147	152

Notes: This table presents structural estimation results for information demand in Experiment 1. The estimation in column (1) is based on a nested logit model, which includes an additional parameter ρ , which is a function of the correlation between the error terms of decisions in the first stage (whether or not to open the envelope) and that of decisions in the second stage (donation, conditional on opening). The estimated ρ (and standard deviation) is 0.52 (0.07). Given the lack of variation in decisions after opening the envelope in the Self treatments, the estimates presented are based on a multinomial logit model. Robust standard errors, clustered at the individual level, are shown in parentheses.

The estimated risk aversion (CRRA parameter) in the MEG treatment is 1.02

($sd=0.02$). The degree of risk aversion in this treatment is thus not significantly different from that in Self-5, where it is 0.97 (t -test, $p=0.2221$), as shown in column (2) of Table 3, which is in line with the benchmarking exercise above.¹⁵

4.1.2 Behavior Conditional on Information Choices

In the MEG treatment, individuals who choose ignorance pick the selfish payment in a large majority of cases (88.4%, on average). Individuals who choose to obtain information exhibit a higher likelihood of giving. The share of individuals who donate when they learn the envelope contains a \$10 donation increases from 49% to 100%, as the price of information increases.¹⁶ This finding is consistent with selection, and with the theoretical framework, as we detail in section 4.2.

In the Self treatments, conditional on demanding information, individuals choose the envelope when it is full 99.2% and 99.0% of the time, in Self-5 and Self-10, respectively. When it is empty, they choose the outside payment 98.0% and 99.4% of the time, respectively. The latter findings confirm that subjects paid attention to their decisions.¹⁷

4.2 Moral Ignorance and Willingness to Donate

Next, we test the relationship between the individual's willingness to donate (WTD) and her demand for information. First, we find a significantly negative relationship between an individual's monetary equivalent of the \$10 donation (or WTD) and her willingness to pay for information, as shown in Figure 6(a). The Spearman correlation coefficient is -0.39 ($p < 0.01$). This finding is consistent with Hypothesis 2.

We also classify our population into five subgroups, depending on the monetary valuation of the donation, to explore the demand for information. Figure 6(b) shows the information-demand curve for each subgroup. The darker curves indicate lower

¹⁵Because little variation occurs in decisions once the envelope is open (less than 2% of individuals choose a dominated option) in the Self treatments, we estimate a simple multinomial logit model. The results of this estimation are shown in columns (2) and (3) of Table 3. Also, we note the CRRA estimates obtained based on decisions in the envelope game in the Self-5 and Self-10 treatments are very similar to those obtained when estimating CRRA coefficients from the control measures after the envelope game, which are 0.77 ($s.d. = 0.02$) in the Self-10 treatment and 1.02 ($s.d. = 0.02$) in the Self-5 treatment.

¹⁶Detailed information on the distribution of information choices by price of information, as well as behavior conditional on obtaining information, is provided in Appendix D.

¹⁷Individuals who do not demand information choose the \$2.50 payment in 84.2% of the cases in Self-5 and 63.7% of the cases in Self-10.

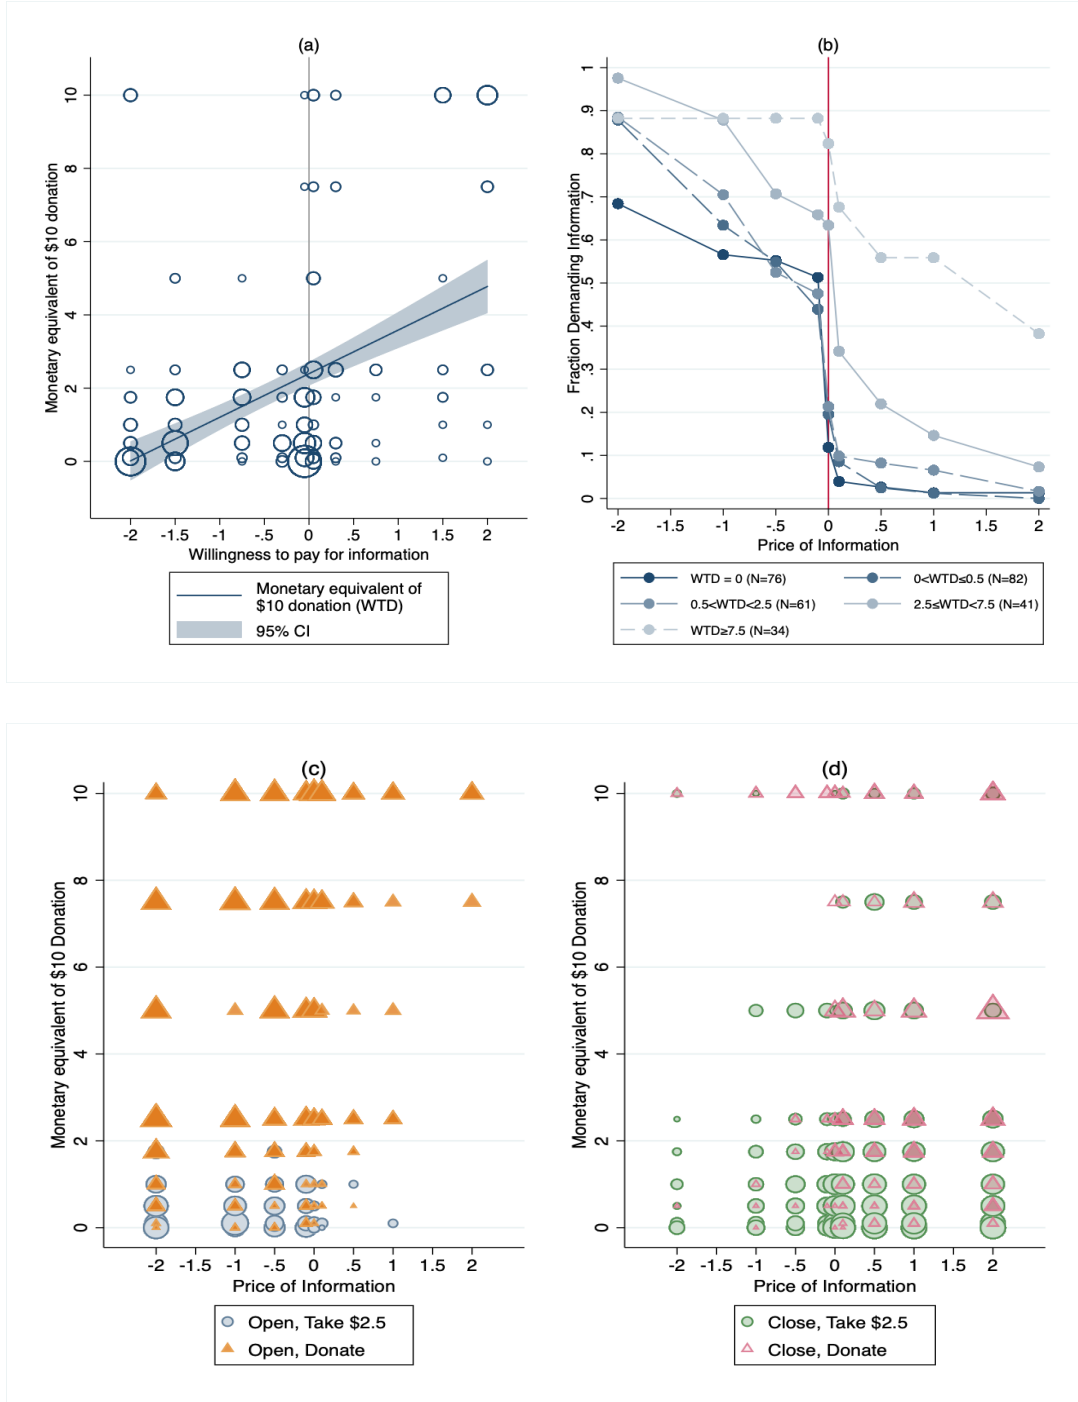


Figure 6: Information Preferences and Donation Preferences

monetary valuations for the donation, whereas the lighter curves indicate higher valuations. In line with the theoretical framework, individuals who value the \$10 donation less than the selfish option ($WTD < 2.5$) often prefer to remain ignorant, even if doing so is costly.

By contrast, individuals who value the \$10 donation more than the selfish option (75 out of 294) often choose to obtain information, even if it is costly. According to Observation 3, individuals with a WTD of \$10 should behave as in Self-10, and those with a WTD of \$5 should act like individuals in Self-5. Comparing those subjects who display a WTD of 10 ($N=26$) in the MEG treatment and subjects in the Self-10 treatment, we find no difference in willingness to pay for information, which is \$0.83 in both cases ($p=0.9773$). Comparing subjects with a WTD of \$5 to the Self-5 does not lead to any significant differences either, yet the number of subjects with a WTD of \$5 in the MEG treatment is small ($N=8$). In Appendix D, we provide further detailed analysis of the behavior of individuals with a WTD above \$2.50.

The theoretical framework also makes predictions regarding individuals' decisions to take the \$2.50 payment or donate. To examine whether behavior in the experiment is consistent with the model, we first examine how often individuals open and take the private payment or open and donate, conditional on finding a donation in the envelope, in Figure 6(c). In line with the theoretical framework individuals with higher WTD choose to open and donate (orange triangles), whereas those with lower WTD choose to open and take the private payment (blue circles). Next, in Figure 6(d), we examine whether individuals choose not to demand information and take the private payment or the envelope with a potential donation. Individuals with high WTD choose to keep the envelope closed and donate more often (red triangles), whereas those with lower WTD choose to keep the envelope closed and take the private payment more often (green circles). Again, these patterns of behavior are broadly consistent with our model.

In addition to the individual's valuation of the \$10 donation, several individual characteristics may explain an individual's willingness to pay for information. In Appendix D, we show the Monitors-Blunters Scale is associated with ignorance, but we find no evidence that the Machiavellianism scale or individual socio-demographic characteristics explain information choices.

4.2.1 Drivers of ignorance

Our theoretical framework proposes that individuals suffer from moral costs when rejecting a certain donation and these costs can drive ignorance. An additional reason for ignorance could be that individuals presented motivated beliefs, which contributed to ignorance. We added a measure of beliefs at the end of the envelope

game, in which we elicited individuals' beliefs regarding the likelihood that the envelope contained a donation for a group of seven other subjects. We incentivized correct answers with a \$0.25 payment. We do not find evidence for self-serving beliefs among those individuals who remain ignorant compared to those who do not.¹⁸

A further question is what is the source of moral discounting, which drives the differential disutility from saying no to a certain donation in the model. Two potential sources are self-image concerns (e.g., Bénabou and Tirole, 2011; Bénabou, Falk and Tirole, 2018) and having a minimal excuse to justify not giving because of uncertainty of consequences (Exley and Kessler, 2018; Gneezy et al., 2017, Engel and Szech 2018, Falk, Neuber, Szech 2019). Keeping the envelope closed can potentially serve as a justification for selfish behavior, because individuals can use it to explain their behavior to others or expect others to use such excuses as well. As argued in Grossman and van der Weele (2017), a repeated choice context may provide little room for self-image concerns, unless subjects engage in narrow bracketing. More than 84% of individuals open the envelope in the first decision of the MEG. After opening, there was a 50 percent chance they would be confronted with the donation opportunity in this first decision. Among those who found a donation in the envelope in the first decision, willingness to pay for avoidance turns out to be somewhat stronger than that of subjects who were not confronted with the donation opportunity yet. Precisely, the willingness to pay for ignorance in those who found a full envelope is \$0.29, compared to a willingness to pay for information of \$0.08 for those who found an empty envelope first. These amounts are significantly different from each other (t -test, $p < 0.01$). Thus, this result suggests that some subjects may have initially underestimated their levels of moral discounting.

4.3 Social Norms Regarding Moral Ignorance

Experiment 1 focused on the impact of monetary incentives on moral ignorance and documented significant elasticity. In Experiment 2 we investigate whether a non-monetary incentive, such as social norms, can affect ignorance. First, we elicited social norms regarding moral ignorance. Figure 7 displays the fraction of individuals who consider taking the \$2.50 payment without opening to be morally appropriate,

¹⁸We find that subjects who pay for ignorance are somewhat more likely to believe that the envelope contains a donation than those who pay for information (Spearman rank correlation coefficient $\rho = -0.20$, $p < 0.001$). This suggests that individuals who avoid did not self-servingly believe that the envelope was empty.

as well as the fraction who considers demanding information by first opening the envelope to be morally appropriate. An action is defined as morally appropriate if an individual considers it somewhat or very morally appropriate.¹⁹




Figure 7: Social Norms regarding Ignorance

On average, less than a third of subjects consider avoiding information and choosing the \$2.50 payment very or somewhat morally appropriate. By contrast, over 70% of individuals consider seeking information very or somewhat morally appropriate. In Appendix D, we provide a detailed analysis of the price sensitivity of norms. Overall, the moral inappropriateness of each action in the envelope game shows relatively little sensitivity to price. Hence, we use the average behavior to provide information on social norms to a new set of participants.

¹⁹Alternatively, we could measure moral-appropriateness ratings by giving a value of -1 to a rating of very morally inappropriate, a value of -1/3 to a rating of somewhat morally inappropriate, a value of 1/3 to a rating of somewhat morally appropriate and a value of 1 to a rating of very morally appropriate (see also, Krupka and Weber, 2013). The results are qualitatively similar. Furthermore, we also elicited the moral appropriateness of choosing the envelope without knowing whether it contains a donation for certain. A large majority of subjects consider such action morally appropriate. Detailed results for all actions are presented in Appendix D.

4.4 The Impact of Social Norms

Our second experiment addresses the question: Can social norms reduce moral ignorance? Figure 8 depicts the demand curve for information when individuals receive social-norm information and when they do not. Avoidance in the NoNorm treatment is shown with the black solid line. Avoidance in Norm-Avoid is depicted by the orange dashed line, and by the orange solid line in Norm-Seek.

Information demand displays a significant drop around \$0 (of 32 percentage points) in the NoNorm treatment, as in Experiment 1. The data thus confirm that a shift from small monetary costs to small monetary rewards for seeking information can affect moral ignorance in a pronounced and robust way.




Figure 8: Information Demand in Experiment 2, by treatment

On average, we do not observe a significant effect of social norms on moral ignorance. Table 4 presents the results from estimating the effect of social norms on ignorance. We observe that ignorance decreases directionally by 1 to 4 percentage points, a change that is not significantly different from zero.

Yet, we observe a positive effect of social norms on WTD. Specifically, the monetary equivalent of the donation increases by \$0.71 (t -test, p -value=0.06) and \$0.68

Table 4: Information Demand in Response to Social Norm Information

<i>Treatment</i>	(1)	(2)	(3)	(4)
	NoNorm	Norm-Avoid	Norm-Seek	All
Price (of Information)	-0.1922*** (0.0151)	-0.1384*** (0.0142)	-0.1351*** (0.0136)	-0.1922*** (0.0151)
Costly Information	-0.3247*** (0.0300)	-0.2941*** (0.0295)	-0.2853*** (0.0286)	-0.3247*** (0.0299)
Costly Information X Price	0.1006*** (0.0247)	0.0616*** (0.0221)	0.0288 (0.0234)	0.1006*** (0.0247)
Norm-Avoid				0.0113 (0.0447)
Norm-Seek				0.0442 (0.0437)
Price X Norm-Avoid				0.0538*** (0.0207)
Price X Norm-Seek				0.0570*** (0.0203)
Costly Information X Norm-Avoid				0.0306 (0.0420)
Costly Information X Norm-Seek				0.0394 (0.0414)
Costly Information X Price X Norm-Avoid				-0.0390 (0.0331)
Costly Information X Price X Norm-Seek				-0.0718** (0.0340)
Constant	0.5796*** (0.0315)	0.5909*** (0.0318)	0.6238*** (0.0304)	0.5796*** (0.0315)
Observations	1,800	1,809	1,872	5,481
R-squared	0.3460	0.2399	0.2579	0.2814
Nr of subjects	200	201	208	609

Notes: This table examines the impact of price on information demand in the NoNorm, Norm-Avoid, and Norm-Seek treatments, using linear probability models. The dependent variable takes a value of 1 if the individual demands information (opens envelope). Robust clustered standard errors are shown in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

(t -test, p -value=0.08) in the Norm-Avoid and Norm-Seek treatments, respectively, compared to the NoNorm treatment. In line with this result, we find that conditional on demanding information, individuals who learn the envelope contains a \$10 donation donate 67% of the time in the NoNorm treatment, 77% in the Norm-Avoid treatment, and 74% in the Norm-Seek treatment. The increase is marginally significant ($p=0.051$ and $p=0.063$, respectively). This in turn increases the likelihood that a donation is made in the MEG with social norms. Compared to the NoNorm treatment, the likelihood increases from 21.7% to 26.2% in the Norm-Avoid treatment

($p=0.052$) and 24.2% in the Norm-Seek treatment ($p=0.267$).²⁰

In terms of our model, we thus observe an increase in the individual’s WTD (i.e., an increase in α). Yet the impact on information demand is small. As expected, selfish individuals seem to have stronger preferences for ignorance than before.

Therefore, we explore the effects of social norms depending on whether information is costly or not. The results are displayed in column (4) of Table 4. Price sensitivity decreases significantly, by approximately 5 percentage points, when ignorance is costly, in both Norm treatments. Thus, a higher share of selfish individuals foregoes sizable monetary payments in order to remain ignorant. By contrast, when obtaining information is costly, subjects become less price sensitive through norm provision (significantly so in the Norm-Seek treatment). For this range of prices, a directional upward shift in the demand curve occurs, as altruistic individuals demand information more often, even if it is costly.

Consistent with these findings, when we structurally estimate our theoretical model in exploratory analyses, we find that individuals’ average altruism $\hat{\alpha}$ increases from 0.21 in the absence of norm information to 0.24 and 0.22 in the Norm-Avoid and Norm-Seek treatments ($p=0.05$ and 0.61, respectively), as shown in Table 5. The estimated moral cost of rejecting a certain donation, $\hat{\beta}$, is 0.84 without norm information in Experiment 2, and it decreases directionally to 0.82 and 0.79 in the Norm-Avoid and Norm-Seek treatments (p -value=0.59 and 0.26, respectively).

Taken together, we find that social norms increase WTD, in line with Hypothesis 3, but do not increase information demand overall. These effects suggest social norms not only increase individuals’ valuation of the donation, but also the moral cost of rejecting a donation, leading to partially offsetting effects of social-norm interventions in still-selfish subjects.

5 Persistence of Ignorance across Contexts

We have documented that information demand about a moral dilemma, such as that studied in the MEG, is highly elastic to monetary incentives but rather inelastic to social norms. An important question is whether information demand in the MEG is predictive of information demand in other moral dilemmas. If so, individual heterogeneity in the willingness to obtain information about moral dilemmas could

²⁰Detailed results regarding donation decisions when the envelope contains a donation and overall likelihood of a donation, by price of information, are shown in Appendix D.

Table 5: Estimation of Preferences for Ignorance with Norms Information

	(1) NoNorm	(2) Norm-Avoid	(3) Norm-Seek
Altruism parameter $\hat{\alpha}$	0.2156 (0.0076)	0.2355 (0.0068)	0.2213 (0.0082)
Moral cost parameter $\hat{\beta}$	0.8379 (0.0253)	0.8165 (0.0311)	0.7924 (0.0325)
Risk aversion parameter (CRRA) \hat{r}	0.9869 (0.0312)	0.9188 (0.0317)	0.9510 (0.0331)
Loglikelihood	-2461.32	-2546.35	-2630.77
Observations	1,800	1,809	1,872
Nr of subjects	200	201	208

Notes: This table presents structural estimation results for information and donation decisions in the Experiment 2. The estimation is based on a nested logit model, which includes an additional parameter ρ , which is a function of the correlation between the error terms of decisions in the first stage (whether to open or not the envelope) and that of decisions in the second stage (donation, conditional on opening). The estimated ρ (and standard deviation) is 0.68 (0.06) in the NoNorm treatment, 0.57 (0.07) in the Norm-Avoid and 0.69 (0.08) in the Norm-Seek treatment in Experiment 2. Robust standard errors, clustered at the individual level, are shown in parentheses.

be seen as a persistent individual trait..

We examine information demand about animal welfare in dairy production, 7 to 10 days after individuals participated in the morally relevant envelope game. We find that a majority of participants (65%) choose to watch the informational video about cows' living conditions, and those watching the video answer the questions correctly in 74% of the cases, compared to 24% when they do not watch it. The main question is whether individuals' willingness to pay for information in the MEG predicts future demand for information in a different moral context, animal welfare. Figure 9 shows a positive relationship between the likelihood of watching the video and an individual's willingness to pay for information in the envelope game.

Table 6 shows the results of a regression analysis of the relationship between willingness to pay to remain ignorant in the MEG and the choice to watch the informational video. As can be seen, a \$1 increase in the individual's willingness to pay for information increases the likelihood that she watches the video about cows' living conditions by 5 percentage points.²¹ This relationship is not driven by an individual's altruism, as the regressions control for the individual's WTD and

²¹Qualitatively similar results are obtained if we study the relationship between the share of correct answers to the questions about the video and ignorance in the donation setting.




Figure 9: Demand for Information across Tasks

this variable has no predictive power on information demand regarding cows' living conditions.

The rate of avoidance of the video varies depending on the social norms individuals were exposed to in Experiment 2. In the NoNorm treatment, the share of individuals who watch the video is 76.9%. It is 68.8% in the Norm-Avoid treatment, and 74.3% in the Norm-Seek treatment. As shown in column (2) of Table 6, the Norm-Avoid treatment led to an increase in avoidance of the video. Although the effect is comparatively small and exploratory, it suggests social-norm interventions should carefully measure short-run as well as long-run impacts, to fully capture potential spillovers onto information demand in other moral contexts.

6 Conclusion

This paper investigates the elasticity of the demand for moral ignorance. Ignorance often enables individuals to engage in questionable ethical decisions in a variety of domains. Yet, little is known about how ignorance can be changed, using monetary and non-monetary incentives.

Table 6: Persistence of ignorance across contexts

	Likelihood of watching video about cows' living conditions	
Willingness to pay for information	0.0509*** (0.0169)	0.0500*** (0.0168)
Monetary equivalent of \$10 donation (WTD)	0.0076 (0.0056)	0.0087 (0.0056)
Norm-Avoid Treatment		-0.0851* (0.0457)
Norm-Seek Treatment		-0.0460 (0.0465)
Observations	533	533

Notes: This table reports marginal effects from probit regressions on the likelihood of watching the informational video about cows' living conditions. The variable willingness to pay for information is measured by the price of information at the point at which the individual switches from demanding information to avoiding information. The variable monetary equivalent of a \$10 donation (WTD) is the individual's willingness to donate as measured after the envelope game. Norm-Avoid and Norm-Seek treatment are treatment dummies for the corresponding treatment in Experiment 2. All regressions include controls for the participant's gender, age, and educational achievement. Clustered standard errors are shown in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

In this paper we show that moral ignorance is highly elastic to monetary incentives. Removing any (small) monetary costs of information and introducing small monetary incentives for information seeking can reduce moral ignorance by roughly 50%. By contrast, larger monetary incentives have comparatively little impact on reducing moral ignorance and come at a much larger cost.

In policy circles, costless interventions such as moral “nudges” are often discussed as interventions to reduce unethical behavior. Our findings with two different framings of social norms suggest social norms perform differently when applied to encouraging information demand, compared to encouraging ethical behavior. Moral nudges in the form of social norms can increase moral behavior. Yet, their impact on curbing moral ignorance is limited. Possibly, norm information increases the moral costs from rejecting moral outcomes and thereby fosters ignorance in rather selfish subjects.

The elasticity of moral ignorance is therefore unlike the elasticity of morally-neutral information demand and the elasticity of moral behavior. It is much more sensitive to small monetary incentives than morally-neutral information demand. At the same time, it is less responsive to social norms, compared to moral behavior.

Understanding the drivers of moral ignorance in the context of a donation de-

cision, as we do in the main part of this paper, not only provides insights about information demand in the context of charitable decisions. Our follow-up experiment shows that ignorance is persistent across time and onto other moral decisions, such as animal welfare information, which suggests that the findings could provide insights relevant to various domains of moral ignorance.

Online Appendix

A Theoretical Analysis

A.1 Proofs

Proof of Proposition 1. We solve the MEG by backwards induction, using the game tree depicted in Figure 1. We begin with the decision to take or donate after an initial decision to open. In this case, taking the selfish amount if $u(\beta \cdot 2.5) > u(\alpha \cdot 10)$ is strictly optimal. This condition is equivalent to $\alpha < \beta \cdot \frac{1}{4}$, due to the monotonicity of u . At the other endnode, after an initial decision not to open, taking the selfish monetary amount if $u(2.5) > \frac{1}{2} \cdot u(\alpha \cdot 10)$ is optimal. By our assumption that $u(x) = x^r$, $r > 0$, this condition is equivalent to $\alpha < 2^{\frac{1}{r}} \cdot \frac{1}{4}$. We can thus turn to the initial decision to open the envelope and its dependence on α . The utility from opening is²²

$$\frac{1}{2}u(2.5) + \frac{1}{2}u(\beta \cdot 2.5) \quad \text{if } \alpha < \beta \cdot \frac{1}{4}$$

and

$$\frac{1}{2}u(2.5) + \frac{1}{2}u(\alpha \cdot 10) \quad \text{if } \alpha \geq \beta \cdot \frac{1}{4}.$$

The utility from not opening is

$$u(2.5) \quad \text{if } \alpha < 2^{\frac{1}{r}} \cdot \frac{1}{4}$$

and

$$\frac{1}{2}u(\alpha \cdot 10) \quad \text{if } \alpha \geq 2^{\frac{1}{r}} \cdot \frac{1}{4}.$$

As $\beta < 1 < 2^{\frac{1}{r}}$, we thus distinguish three cases depending on the location of α . If α lies below both thresholds, $\alpha < \frac{\beta}{4}$, donating is suboptimal regardless of the decision

²²Here and in the following, we assume an agent who is indifferent between taking the money and donating will donate. Similarly, the agent favors options with a higher donation probability in case of indifference.

in the first stage. In this case, the comparison

$$\frac{1}{2} \cdot u(2.5) + \frac{1}{2} u(\beta \cdot 2.5) < u(2.5) \quad (1)$$

implies leaving the envelope closed is optimal.²³ In the intermediate case when $\beta \cdot \frac{1}{4} \leq \alpha < 2^{\frac{1}{r}} \cdot \frac{1}{4}$, we have to compare the utility of $\frac{1}{2} \cdot u(2.5) + \frac{1}{2} \cdot u(\alpha \cdot 10)$ from opening and $u(2.5)$ from leaving the envelope closed. Opening is thus optimal for $\alpha \geq \frac{1}{4}$ whereas leaving the envelope closed is optimal otherwise. In the third case $\alpha \geq 2^{\frac{1}{r}} \cdot \frac{1}{4}$, the relevant comparison is between $\frac{1}{2} \cdot u(2.5) + \frac{1}{2} \cdot u(\alpha \cdot 10)$ and $\frac{1}{2} \cdot u(\alpha \cdot 10)$. In this case, opening the envelope is optimal. ■

Proof of Proposition 2. The behavior at the endnodes is not affected by the additional costs of opening or leaving the envelope closed. After opening and finding a full envelope, the agent donates if $\alpha \geq \frac{\beta}{4}$. If the envelope is kept closed, the agent donates if $\alpha \leq \frac{1}{2}$.

For the initial opening decision, we distinguish between three cases, depending on whether $\alpha < \frac{\beta}{4}$, $\alpha \in [\frac{\beta}{4}, \frac{1}{2})$, or $\alpha \geq \frac{1}{2}$. First, for $\alpha < \frac{\beta}{4}$, the relevant comparison is now between a utility of $\frac{1+\beta}{2} 2.5 + m_o$ from opening and $2.5 + m_c$ from keeping it closed. Opening is strictly dominant if the difference between m_o and m_c is positive and sufficiently large,

$$m_o - m_c > 5 \left(\frac{1}{4} - \frac{\beta}{4} \right).$$

Second, for $\alpha \in [\frac{\beta}{4}, \frac{1}{2})$, the comparison is between a utility of $\frac{2.5+\alpha \cdot 10}{2} + m_o$ from opening and $2.5 + m_c$ from keeping the envelope closed. Opening strictly dominates if

$$m_o - m_c > 5 \left(\frac{1}{4} - \alpha \right). \quad (2)$$

Otherwise, leaving the envelope closed is best. Observe the right-hand side of (2) switches signs at $\alpha = \frac{1}{4}$. Thus, if $\alpha < \frac{1}{4}$, a positive value of $m_o - m_c$ is needed to motivate the agent to open the envelope. By contrast, for $\alpha > \frac{1}{4}$, the agent will still open the envelope when m_c is slightly larger than m_o .

In the third case $\alpha \geq \frac{1}{2}$, that is, for subjects with a very high valuation for the donation, we have to compare $\frac{2.5+\alpha \cdot 10}{2} + m_o$ from opening and $\frac{\alpha \cdot 10}{2} + m_c$ from keeping

²³In the boundary case $\beta = 1$, the agent is instead indifferent between opening and not opening. This case is the only part of the analysis that changes for $\beta = 1$.

the envelope closed. Opening strictly dominates if

$$m_o - m_c > -\frac{5}{4};$$

that is, unless m_c is quite high, opening the envelope is best. ■

A.2 Cases of Risk Aversion and Risk Lovingness

Individuals may have different risk attitudes. Intuitively, risk aversion makes the closed envelope less attractive such that even under high monetary incentives, only very altruistic subjects prefer the closed envelope. The following figure demonstrates the case of $u(x) = \sqrt{x}$ for different levels of moral discounting, β . If the moral discounting is pronounced, even most selfish individuals prefer to leave the envelope closed in order to avoid moral costs from rejecting the donation.

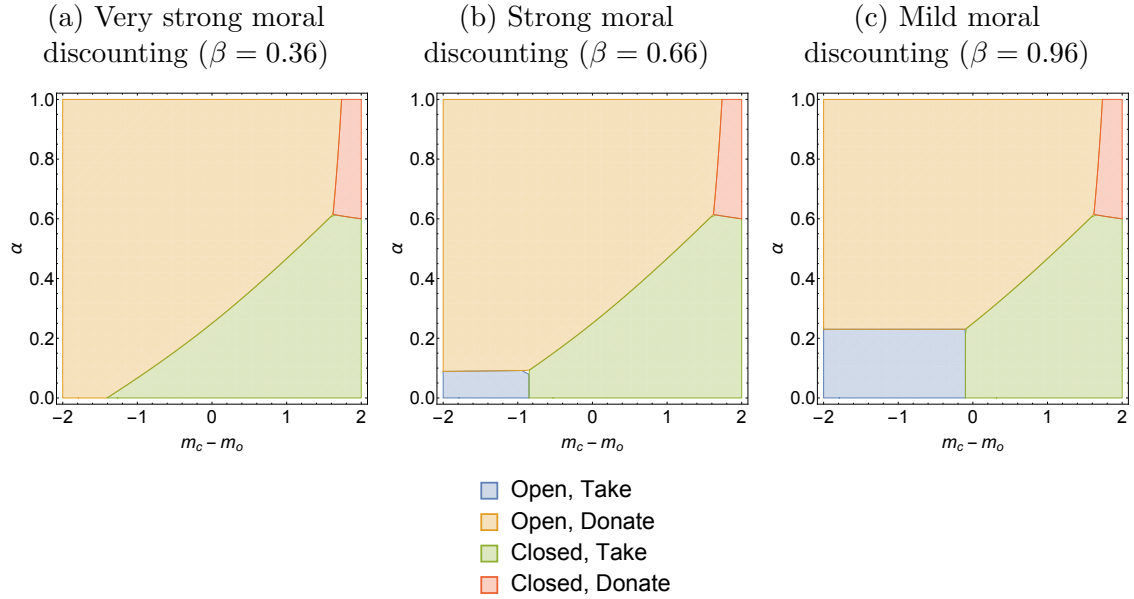


Figure A.1: Risk Aversion: $u(x) = \sqrt{x}$; $\beta = 0.36, 0.66, 0.96$, respectively

By contrast, the closed envelope can become quite appealing for altruists if they are risk loving. The following figure illustrates that case.

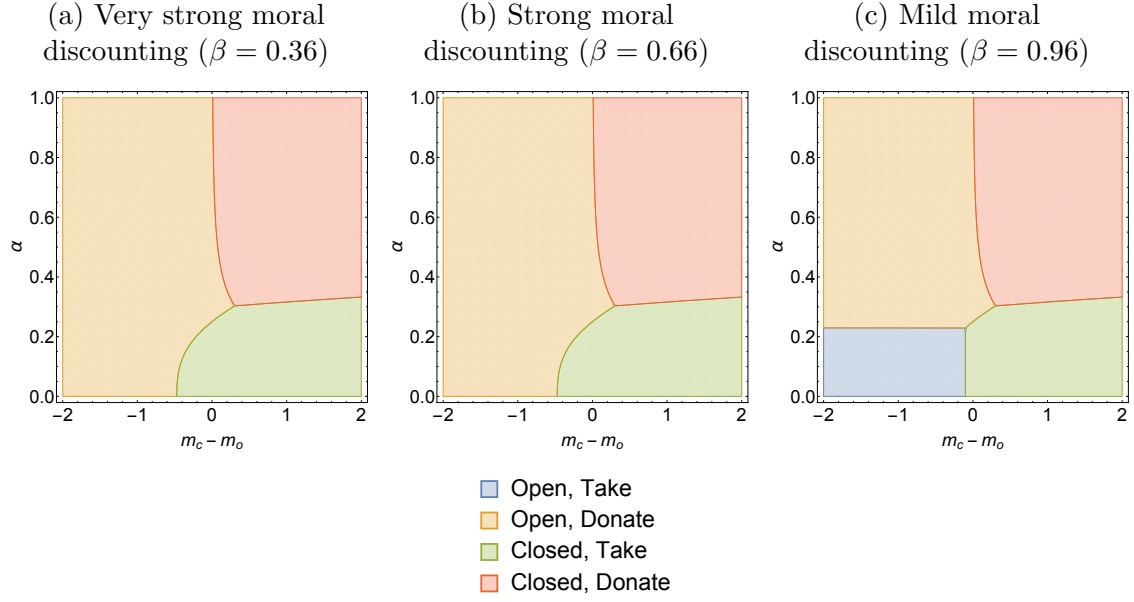


Figure A.2: Risk Loving: $u(x) = x^4$; $\beta = 0.36, 0.66, 0.96$, respectively

B Instructions

Below, we present the instructions for the MEG treatment in Experiment 1. The Self-5 and Self-10 treatments had the same instructions except that the \$10 donation was replaced by a \$5 or \$10 payment for the individual. In Experiment 2, we added information on social norms at the end of the instructions, as indicated in brackets below.

In this study, you make decisions involving money for you and a donation to the Malaria Consortium in your name.

Your donation takes place via **an envelope**.

The envelope either contains: a \$10 donation on your behalf with 50% chance, or no donation with 50% chance

page break

In this part, you **do not know what the envelope contains**. You decide whether to get \$2.50, get the envelope, or reveal what the envelope contains first.

You will make 9 decisions. In each decision you have **three options**:

- (a) **Get \$2.50:** then, you get \$2.50.
- (b) **Get the envelope:** then, you donate what the envelope contains, which you do not know.
- (c) **Reveal what the envelope contains first:** then, you are shown whether the envelope contains **a \$10 donation on your behalf** or **no donation**. After being informed of the envelope's content, you decide either to get \$2.50 or get the envelope.

In each decision, you may receive **an additional amount for choosing option (c) 'Reveal what the envelope contains first', or you may receive an additional amount for NOT choosing option (c)**, that is not 'revealing what the envelope contains first', and choosing options (a) or (b).

Across the 9 decisions, this additional amount you receive varies, from \$2 for revealing what the envelope contains to \$0, and from \$0 to \$2 for NOT revealing what the envelope contains.

You will not know which is the 'decision that counts' until the end of the study. Because the computer is making a random draw, any of the choices could be the 'decision that counts.' Therefore, you should think carefully about the choice you make in each question.

page break

Remember, the envelope either contains:

- a \$10 donation** on your behalf with 50% chance, or
- no donation with 50% chance.**

In what follows you will be shown an example and will be asked to answer several questions, before making your decisions.

page break

EXAMPLE

As an example, let us consider question 5. In this question, you receive **\$0 for revealing what the envelope contains, and \$0 for not revealing what the envelope contains.** The question is shown below.

5. If you get \$0 for revealing and \$0 for not revealing what the envelope contains first, what do you choose?

- (a) Get \$2.50
- (b) Get the envelope
- (c) Reveal what the envelope contains

If you choose '(a) get \$2.50', **you get \$2.50.**

If you choose '(b) get the envelope', **you donate \$10 with 50% chance or donate \$0 with 50% chance.**

If you choose '(c) reveal what the envelope contains', you learn what is inside the envelope. Two cases can then happen:

1. **The envelope contains a \$10 donation on your behalf.** Then you choose between:

Get \$2.50 or

Donate \$10.

2. **The envelope contains no donation.** Then you choose between:

Get \$2.50 or

Donate \$0.

Thus, choosing (c) brings you to another choice. This choice is either between a monetary amount for you and a donation on your behalf or between a monetary amount for you and no donation in your name.

page break

CONTROL QUESTIONS

As an exercise, let us consider question 1.

The envelope either contains a \$10 donation with 50% chance or no donation with 50% chance.

In question 1, you **receive \$2 for revealing what the envelope contains.** You choose from these options:

- (a) Get \$2.50.
- (b) Get the envelope.
- (c) Reveal what the envelope contains (plus \$2 for revealing what the envelope contains).

Suppose you choose (a). What happens?

- ☐ I get \$0
- ☐ I get \$2.50

Suppose you choose (b). What happens?

☐ I get \$4.50

☐ I donate what is in the envelope.

Suppose you choose (c). What happens?

☐ I learn what is inside the envelope and then decide between the envelope and \$2.50. I either decide between a \$10 donation on my behalf or \$2.50 for me, or I decide between a \$0 donation or \$2.50 for me. I also get \$2 for choosing to reveal what is in the envelope.

☐ I do not learn what is inside the envelope and get \$0.

page break

Now you will make your decisions for this part.

In each question, the chance that the envelope contains a \$10 donation is 50%. So, the envelope's content varies for each question according to chance.

One of your decisions may be the 'decision that counts'. So please decide carefully!

page break

[Treatment Norm-SeekInfo: **NOTE:** More than 70% of MTurkers who evaluated the actions in this part consider it **MORALLY APPROPRIATE** to choose the option '**Reveal what the envelope contains**' first.]

[Treatment Norm-AvoidInfo: **NOTE:** More than 70% of MTurkers who evaluated the actions in this part consider it **MORALLY INAPPROPRIATE** to choose the option '**Get \$2.5**' without revealing what the envelope contains first.]

C Elicitation of Control Measures

After the main part of the experiment, we elicited the monetary equivalent of a certain \$10 donation, by asking the individual to make eight binary choices between the donation and payments to her that increased from \$0.10 to \$10. Each choice between a private payment and the donation was presented in a separate screen, and across screens the value of the private payment increased. Thereafter, we elicited the monetary equivalent of a \$10 donation that occurs with a 50% chance. Individuals again made eight binary choices, each between the potential donation and a payment that increased from \$0.01 to \$5. These choices were elicited in the MEG treatment in Experiment 1 and all treatments in Experiment 2.

For each individual, we calculate her monetary equivalent of a certain (uncertain) donation as the maximum value of the payment to her that she was willing to give up instead of the donation. As shown in Table B.1., on average, individuals' monetary equivalent of a certain \$10 donation was 1.91 ($sd = 2.94$), whereas it was 0.69 ($sd = 0.89$) for a 50% chance of a \$10 donation in Experiment 1.

Table C.1: Control Measures and Sample Characteristics

		Experiment 1				Experiment 2	
		Donation	Self-5	Self-10	NoNorm	Norm-Avoid	Norm-Seek
Monetary equivalent of:							
(1) \$10 Donation	Mean	1.91	-	-	2.92	3.60	3.56
	SD	2.94			3.48	3.78	3.92
(2) \$10 Donation/Self payment, with $p = 0.5$	Mean	0.69	1.40	2.59	0.78	0.81	0.72
	SD	0.82	0.76	1.83	0.88	0.97	0.89
Subject characteristics							
Female	Mean	45.9%	42.9%	46.7%	55.0%	51.2%	52.9%
Age	Mean	36.5	37.6	35.3	36.3	35.6	38.0
High school graduate	Mean	40.5%	30.6%	44.1%	36.0%	31.8%	33.2%
On Mturk 7 days a week	Mean	92.5%	91.8%	90.8%	89.0%	86.1%	91.8%

In the Self 5 and Self 10 treatments of Experiment 1, we elicited the certainty equivalent of a \$5 and \$10 payment that occurred with a 50% chance. We asked the individual to make eight binary choices between the uncertain payment and

payments to her that increased from \$0.50 to \$5 in the Self 5 treatment, and \$1 and \$10 in the Self 10 treatment. On average, the certainty equivalent of a 50% chance of \$5 was 1.40, and that of a 50% chance of \$10 was 2.59.

The second part of Table B.1. displays the characteristics of subjects who participated in Experiment 1 and 2, including gender, age, high school graduates, and intensity of work at Amazon Mechanical Turk.

C.1 Details of MTurk assignments

We used TurkPrime to invite individuals to participate in our studies on MTurk, and re-invite them for the follow-up task. In Experiments 1 and 2, individuals always received a fix payment of \$3, in addition to their earnings from the experiment. The criteria to participate were that they should be located in the US and have an approval rate of at least 80%. The median time to complete the study was approximately 20 minutes, and more than 96% of participants who started completed the experiment.

D Additional Results

D.1 Distribution of information choices in Experiments 1 and 2

Table D.1 below presents the distribution of choices in Experiment 1. For each price of avoidance, we show the percentage of individuals who (a) avoid and choose \$2.50 (“Choose \$2.50”), (b) avoid and donate (“Choose envelope”), (c) seek information (“Open envelope”).

Table D.1: Distribution of Choices in Experiment 1

Price of information		Treatment		
		Donation	Self-5	Self-10
-\$2	Choose \$2.5	14.6%	7.5%	2.6%
	Choose envelope	1.0%	2.0%	2.6%
	Open envelope	84.4%	90.5%	94.7%
- \$1	Choose \$2.5	28.6%	18.4%	6.6%
	Choose envelope	2.0%	2.0%	3.3%
	Open envelope	69.4%	79.6%	90.1%
- \$0.50	Choose \$2.5	36.7%	21.8%	10.5%
	Choose envelope	2.7%	4.1%	4.6%
	Open envelope	60.5%	74.1%	84.9%
- \$0.10	Choose \$2.5	42.5%	23.1%	13.8%
	Choose envelope	2.7%	5.4%	5.9%
	Open envelope	54.8%	71.4%	80.3%
\$0	Choose \$2.5	62.6%	25.9%	15.1%
	Choose envelope	6.1%	8.2%	9.9%
	Open envelope	31.3%	66.0%	75.0%
\$0.10	Choose \$2.5	72.8%	34.0%	19.1%
	Choose envelope	9.2%	6.1%	11.2%
	Open envelope	18.0%	59.9%	69.7%
\$0.50	Choose \$2.5	77.2%	42.2%	21.7%
	Choose envelope	10.2%	7.5%	13.2%
	Open envelope	12.6%	50.3%	65.1%
\$1	Choose \$2.5	76.5%	59.9%	30.9%
	Choose envelope	12.9%	9.5%	18.4%
	Open envelope	10.5%	30.6%	50.7%
\$2	Choose \$2.5	76.5%	76.2%	42.1%
	Choose envelope	17.3%	12.9%	23.7%
	Open envelope	6.1%	10.9%	34.2%

Table D.2 below presents the distribution of choices in Experiment 2.

Table D.2: Distribution of Choices in Experiment 2

Price of information		Treatment		
		NoNorm	NormAvoid	NormSeek
-\$2	Choose \$2.5	5.0%	12.4%	9.1%
	Choose envelope	2.0%	4.0%	4.8%
	Open envelope	93.0%	83.6%	86.1%
-\$1	Choose \$2.5	16.5%	19.4%	15.9%
	Choose envelope	4.0%	4.5%	5.3%
	Open envelope	79.5%	76.1%	78.8%
-\$0.50	Choose \$2.5	19.5%	21.4%	19.2%
	Choose envelope	5.5%	7.5%	5.3%
	Open envelope	75.0%	71.1%	75.5%
-\$0.10	Choose \$2.5	26.0%	24.9%	24.5%
	Choose envelope	6.5%	8.5%	5.8%
	Open envelope	67.5%	66.7%	69.7%
\$0	Choose \$2.5	49.0%	42.8%	41.3%
	Choose envelope	7.0%	9.5%	8.2%
	Open envelope	44.0%	47.8%	50.5%
\$0.10	Choose \$2.5	58.5%	53.7%	51.0%
	Choose envelope	15.5%	16.4%	14.9%
	Open envelope	26.0%	29.9%	34.1%
\$0.50	Choose \$2.5	61.5%	53.7%	52.9%
	Choose envelope	17.5%	19.9%	18.3%
	Open envelope	21.0%	26.4%	28.8%
\$1	Choose \$2.5	64.5%	56.7%	60.1%
	Choose envelope	22.0%	23.9%	19.7%
	Open envelope	13.5%	19.4%	20.2%
\$2	Choose \$2.5	64.5%	57.7%	61.5%
	Choose envelope	27.0%	26.9%	24.5%
	Open envelope	8.5%	15.4%	13.9%

D.2 Illustration of calibration in Experiment 1

Figure D.1 illustrates the equivalence between the MEG treatments and the Self-5 and Self-10 treatments, as discussed in the main text. The black connected line shows the willingness to pay for information in the Self treatments. The red line indicates the willingness to pay for ignorance in the MEG treatment.




Figure D.1: Benchmarking Ignorance in MEG treatment to Information Demand in Self Treatments

D.3 Behavior conditional on information demand in Experiments 1 and 2

Individuals who choose to seek information learn whether the envelope actually contains a donation. If subjects learn the envelope is empty, they choose the \$2.50 payment to themselves in 93.5% of the cases. If subjects learn the envelope contains a \$10 donation, between 49% and 100% choose the donation, as shown in Table D.3. Consistent with selection of those who value the donation opportunity highly into

seeking information, the share of those who donate increases as the cost of ignorance decreases.

Table D.3: Behavior Conditional on Seeking Information in Experiment 1

Price of information	(a)	N	(b)	N
	Envelope contains \$10 Donation % choose envelope		Envelope contains nothing % choose envelope	
-\$2	49.2%	126	5.7%	122
-\$1	50.5%	101	1.9%	103
-\$0.50	52.1%	94	4.8%	84
-\$0.10	51.3%	78	2.4%	83
\$0	74.5%	47	6.7%	45
\$0.10	87.1%	31	4.5%	22
\$0.50	96.0%	25	16.7%	12
\$1	93.3%	15	6.3%	16
\$2	100.0%	7	9.1%	11

Notes: This table shows the percentage of individuals choosing the envelope in two cases: (a) when the envelope contains a \$10 donation and (b) when it is empty. In each case, we also show the number of observations.

Table D.4: Behavior Conditional on Seeking Information and Finding Donation in Experiment 2

Price of avoidance	Percentage choosing envelope if envelope contains \$10 donation					
	Donation-NoNorm	Norm-Avoid		Norm-Seek		
	% choose envelope	N	% choose envelope	N	% choose envelope	N
\$2	56.8%	88	66.7%	81	65.6%	93
\$1	62.3%	77	65.3%	72	70.3%	74
\$0.50	68.1%	72	68.8%	77	67.9%	78
\$0.10	61.8%	68	74.0%	77	67.2%	64
\$0	72.0%	50	80.3%	61	78.2%	55
-\$0.10	90.9%	22	100.0%	31	97.0%	33
-\$0.50	83.3%	18	100.0%	28	85.3%	34
-\$1	91.7%	12	100.0%	18	95.0%	20
-\$2	92.3%	13	100.0%	15	90.9%	11

Notes: This table shows the percentage of individuals choosing the envelope when the envelope contains a \$10 donation in each treatment in Experiment 2. In each case, we also show the number of observations.

D.4 Donation likelihood in Experiment 2

In Figure D.2 we show the likelihood that a donation is made in each treatment of Experiment 2. The likelihood of a donation stems from two decisions. First, the

individual donates after opening the envelope and finding a donation inside. Second, the individual donates without opening the envelope.




Figure D.2: Likelihood of a donation in Experiment 2

D.5 The determinants of information demand in the MEG treatment

In Table D.5, we examine the determinants of information preferences in the MEG treatment. In addition to the individual's willingness to donate, one important preference is the individual's value of a \$10 donation that occurs only with a 50% chance. To measure how much the individual's value drops when uncertainty cannot be removed, we compare the monetary equivalent of a \$10 donation with certainty and the one with a 50% chance. If the equivalent with uncertainty is less than half of the equivalent with certainty, we classify the individual as risk averse (with respect to the donation). We do not observe that the change in the donation valuation when it is uncertain, relative to when it is certain, is related to information demand. Additional characteristics that could explain information decisions are the score on the Monitors-Blunters Scale and the Machiavellianism scale, as well as gender, age, education, and frequency of work on MTurk. We find the Monitors-Blunters Scale

is associated with ignorance, but we do not find evidence that the Machiavellianism scale or individual socio-demographic characteristics explain information choices.

Table D.5: Determinants of Information Demand in the MEG treatment

	(1)	(2)	(3)
	Willingnes to pay for information		
Monetary equivalent of \$10 donation	0.1625*** (0.0208)	0.1653*** (0.0208)	0.1640*** (0.0213)
Risk averse	0.1168 (0.1485)	0.0992 (0.1482)	0.0942 (0.1496)
Monitors-Blunters Scale Score		0.0273** (0.0135)	0.0277** (0.0138)
Mach IV Score		0.0416 (0.1179)	0.0291 (0.1227)
Female			-0.0623 (0.1206)
Age			0.0005 (0.0056)
High school degree or higher			-0.0744 (0.1188)
Works every day on Mturk			0.0479 (0.2232)
Constant	-0.7416*** (0.0715)	-0.9966*** (0.3518)	-0.9651** (0.4721)
Observations	294	294	294
R-squared	0.2014	0.2130	0.2147

Notes: This table examines the determinants of willingness to pay for information in the MEG treatment. The dependent variable takes values from -2 to 2, depending on when the individual chooses to switch from obtaining information to not obtaining information. The monetary equivalent of a \$10 donation is the individual's valuation of the donation. Risk averse is a dummy variable that takes a value of 1 if the individual values a donation opportunity with a 50% chance less than half of her monetary equivalent of a certain donation. Standard errors are shown in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

D.6 Detailed analysis of the demand for moral ignorance in altruistic subjects

When we examine the behavior of subjects with a WTD between 2.5 and 7.5 shown in Figure 6(b), we see a significant drop in information demand to the right of \$0, which we do not observe in the Self-5 treatment.

This drop could be due to two reasons. The first is risk preferences in the donation domain. If individuals are sufficiently altruistic and risk seeking, when information

becomes costly, they could prefer to donate the closed envelope. Through the lens of the model, we should only see this preference for larger costs of information. In the data, though, among individuals who do not demand information as soon as it becomes costly, only about a third of the subjects (31.4%) decide to donate the closed envelope.²⁴

The second reason for the drop to the right of \$0 may be a change in social norms. Though not captured in our model, one could imagine costs of information, even if tiny, provide some excuse for staying ignorant. We collect data on the moral norms regarding ignorance. These norms turn out to be rather inelastic to the price of information. Nevertheless, information demand is considered somewhat less important when information becomes costly. This norm change may contribute to the drop around \$0 for subjects who have a high valuation for the certain donation.

D.7 Norm-elicitation results in Experiment 2

Table D.6 shows individual ratings of moral appropriateness of each action, for each price of avoidance. Panel A focuses on the choice to avoid and choose \$2.50. Panel B focuses on the choice to avoid and donate by choosing the envelope. Panel C focuses on the choice to seek information, by opening the envelope.

²⁴Among the selected sample of individuals who no longer demand information when it costs \$0.10, we find that 30 individuals choose the \$2.50 payment, whereas 9 choose the closed envelope. To examine whether those subjects who switch to the closed envelope are relatively risk loving with respect to the donation, we explore the ratio of their valuation of the donation with certainty, relative to their valuation of the donation with uncertainty. A risk-neutral individual would exhibit a ratio of 2. We find that, on average, the ratio for all subjects is 2.33 (s.d. 2.78). For the subjects who switched, it is 1.74 (s.d. 0.66). Thus, indeed, we find some indication of risk lovingness in these subjects.

Table D.6: Individual Ratings of Moral Appropriateness of Each Action

	Moral Appropriateness Category Rating			
	Very morally inappropriate	Somewhat morally inappropriate	Somewhat morally appropriate	Very morally appropriate
Price of information	PANEL A. Choose \$2.50			
-\$2	38.24	38.24	18.63	4.9
-\$1	39.22	37.25	16.67	6.86
-\$0.50	40.2	30.39	23.53	5.88
-\$0.10	38.24	32.35	18.63	10.78
\$0	36.27	30.39	24.51	8.82
\$0.1	41.18	30.39	19.61	8.82
\$0.5	38.24	34.31	19.61	7.84
\$1	40.2	32.35	17.65	9.8
\$2	43.14	31.37	17.65	7.84
Total	39.43	33.01	19.61	7.95
Price of information	PANEL B. Choose envelope			
-\$2	0.98	3.92	54.9	40.2
-\$1	0	6.86	55.88	37.25
-\$0.50	0	6.86	54.9	38.24
-\$0.10	2.94	5.88	50.98	40.2
\$0	1.96	7.84	47.06	43.14
\$0.1	0	9.8	44.12	46.08
\$0.5	2.94	6.86	48.04	42.16
\$1	0	7.84	50	42.16
\$2	0.98	9.8	41.18	48.04
Total	1.09	7.3	49.67	41.94
Price of information	PANEL C. Open envelope first			
-\$2	1.96	8.82	45.1	44.12
-\$1	0.98	9.8	46.08	43.14
-\$0.50	0.98	8.82	48.04	42.16
-\$0.10	0.98	9.8	42.16	47.06
\$0	1.96	5.88	47.06	45.1
\$0.1	3.92	17.65	43.14	35.29
\$0.5	2.94	10.78	55.88	30.39
\$1	6.86	10.78	49.02	33.33
\$2	7.84	12.75	45.1	34.31
Total	3.16	10.57	46.84	39.43

Table D.7 presents the results of a linear probability model on the evaluation of each action as morally appropriate. The regression models include an indicator variable for costly information, that is, when prices are strictly positive, to allow for a kink around \$0. The regressions also include an interaction term between the indicator for costly information and the price of information, to allow for a different effect of price on moral appropriateness, depending on whether information is costly. The moral appropriateness of demanding information does not vary significantly with the price of information, but it exhibits a kink around a price of \$0. If information is costly, demanding information is 8 percentage points less likely to be considered morally appropriate.²⁵

Table D.7: Moral Appropriateness

	(1)	(2)	(3)
<i>Action:</i>	Get \$2.5	Morally Appropriate Demand Information	All
Price (of Information)	0.0450** (0.0181)	0.0087 (0.0140)	0.0087 (0.0140)
Costly Information	-0.0257 (0.0241)	-0.0827** (0.0346)	-0.0827** (0.0346)
Costly Information X Price	-0.0595** (0.0297)	-0.0170 (0.0243)	-0.0170 (0.0243)
Get \$2.5			-0.5955*** (0.0492)
Price X Get \$2.5			0.0362 (0.0219)
Costly Information X Get \$2.5			0.0570 (0.0458)
Costly Information X Price X Get \$2.5			-0.0425 (0.0366)
Constant	0.3108*** (0.0440)	0.9063*** (0.0233)	0.9063*** (0.0233)
Observations	918	918	1,836
R-squared	0.0033	0.0150	0.3564
Nr. of subjects	102	102	102

Notes: This table examines the impact of price on the likelihood that getting \$2.50 (private payment) and demanding information is considered very or somewhat morally appropriate, using linear probability models. The dependent variable takes a value of 1 if the individual considers getting \$2.50 (private payment) or demanding information very morally appropriate or somewhat morally appropriate. Robust clustered standard errors are shown in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

²⁵This result provides a potential explanation for the drop in information demand when price increases from \$0 to \$0.10, documented in Experiment 1.

In the Norms treatment, we also elicited the moral appropriateness of donating versus acting selfishly. When faced with a certain \$10 donation, 78% of individuals consider it morally inappropriate not to donate. When faced with a 50% chance of a \$10 donation, 66.7% consider it morally inappropriate not to donate. These results are broadly in line with existing research evaluating social norms around sharing decisions (e.g., Krupka and Weber, 2013).

E Structural Estimation

We structurally estimate the parameters of our theoretical model, using the experimental decisions in Experiment 1 and Experiment 2. As in our model, we assume individuals have CRRA utility, with risk-aversion parameter r , such that $u(x) = x^r$. When individuals donate, they value the donation with α . When they choose the selfish option, knowing the envelope contains a certain donation, they suffer from the (additional) moral cost β . Because the decision structure of individuals involves two steps, we estimate a nested logit model, with three branches (for details, see, e.g., Cameron and Trivedi, 2005, Ch. 12.6.2). Two branches are degenerate. First, if the individual decides to take the selfish payment of \$2.50 without opening, we denote the utility by $V^{c,0}$. Second, if she takes the envelope without opening it, we denote the utility as $V^{c,1}$.

The third branch is the choice to open the envelope. Then, knowing whether the envelope is full or empty, the individual decides whether to take the envelope. To specify the likelihood, denote the decision to take the envelope as $d \in \{0, 1\}$. The utility of d , conditional on opening, is $V^{d,f}$, where f indicates whether the envelope is full or empty. The likelihood of opening and taking the envelope is

$$p_{o,d} = p_o \times p_{d|o} = \frac{\exp(\rho I^o)}{\exp(\rho I^o) + \exp(V^{c,1}) + \exp(V^{c,0})} \times \frac{\exp(V^{d,f}/\rho)}{\sum_d \sum_f \exp(V^{d,f}/\rho)},$$

where $I^o = \ln(\sum_d \sum_f \exp(V^{d,f}/\rho))$, which is known as the inclusive value. The likelihood of leaving the envelope closed and taking it is

$$p_{c,1} = \frac{\exp(V^{c,1})}{\exp(\rho I^o) + \exp(V^{c,1}) + \exp(V^{c,0})},$$

and the likelihood of leaving the envelope closed and taking the \$2.50 payment is:

$$p_{c,0} = \frac{\exp(V^{c,0})}{\exp(\rho I^o) + \exp(V^{c,1}) + \exp(V^{c,0})}.$$

In all estimations, we include the payoff of the decision or resulting donation as well as the individual’s show-up fee. This approach avoids negative payoffs in the rare cases where the individual opens the envelope and chooses the empty envelope.

The estimation of the nested logit includes an additional parameter, ρ , which is a function of the correlation between the error term in the decisions in the first stage (whether to open or not the envelope) and the error term in the decisions in the second stage (donation, conditional on opening). This parameter is noted in the footnote of Tables 3 and 5. Because ρ enters multiplicatively in the utility of choosing between the envelope (with a donation or not) and the outside payment of \$2.50, it cannot be separately identified from a Fechner error (see, e.g., von Gaudecker et al., 2011). Hence, we do not explicitly add Fechner errors to the model, and interpret ρ with care.

To further examine the coherence of the estimated risk-aversion parameters in the Self treatments, we also estimated the implied CRRA parameters from the control measures, elicited through simple binary decisions, after the main part of the experiment (for a more detailed description see Appendix C). The estimated CRRA parameter from those decisions in the Self-10 treatment is 0.77 (sd=0.02), and that in the Self-5 treatment is 1.02 (sd=0.02). Hence, these parameters are closely in line with those estimated from the decisions in the main part of the experiment.

Finally, we note that we have explored the results of structural estimation when including all participants, as well as those who were inconsistent in their decisions. We find the results remain qualitatively similar, and that the effects on social-norm information are strengthened (leading again to an increase in the altruism parameter and an increase in moral costs, measured as smaller β).

E.1 Additional analyses preregistered for Experiment 2

For Experiment 2, we planned to examine the relationship between willingness to pay for information and individual preferences for giving, individual characteristics, and two personality scales related to ignorance and moral behavior. Because the insights from these analyses are very similar to those obtained from Experiment 1, we show these results below instead of in the main text. We did not collect

information on individual preferences for giving in the Norm treatment and hence do not conduct the analysis for that treatment.

Table E.1: Determinants of Information Demand in Experiment 2

	(1)	(2)	(3)
<i>Treatment:</i>	NoNorm	Norm-Avoid	Norm-Seek
	Willingness to pay for information		
Monetary equivalent of \$10 donation	0.0574** (0.0221)	0.1186*** (0.0232)	0.1156*** (0.0207)
Risk averse	0.1626 (0.1728)	0.0686 (0.1868)	0.2477 (0.1959)
Monitors-Blunters Scale Score	0.0330* (0.0171)	0.0457** (0.0191)	0.0166 (0.0169)
Mach IV Score	-0.3003 (0.1880)	-0.1821 (0.1971)	-0.1797 (0.1685)
Female	-0.2096 (0.1487)	-0.1742 (0.1762)	-0.0860 (0.1628)
Age	0.0030 (0.0065)	0.0003 (0.0075)	-0.0139** (0.0065)
High school degree or higher	-0.0869 (0.1523)	0.1206 (0.1832)	-0.3525** (0.1647)
Works every day on Mturk	-0.2457 (0.2349)	-0.4753* (0.2472)	0.0097 (0.2839)
Constant	0.6500 (0.7130)	0.2527 (0.6896)	0.6705 (0.6588)
Observations	200	201	208
R-squared	0.1049	0.1865	0.2135

Notes: This table examines the determinants of willingness to pay for information in the MEG treatment. The dependent variable takes values from -2 to 2, depending on when the individual chooses to switch from obtaining information to not obtaining information. The monetary equivalent of a \$10 donation is the individual's valuation of the donation. Risk averse is a dummy variable that takes a value of 1 if the individual values a donation opportunity with a 50% chance less than half of her monetary equivalent of a certain donation. Standard errors are shown in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

We also planned to examine the relationship between moral views and desire to obtain information in the follow-up task, in which subjects had the opportunity to answer questions about the living conditions of cows. We find that 83% of participants complete the video task (84 out of 101). Of these, 73% watch the video.

We build an index measuring how often an individual considers it morally appropriate to take \$2.50 without opening first (28% of the time), to take the envelope without opening it first (91% of the time), or to open the envelope first (86% of the time).

We find that individuals who consider taking \$2.50 to be more morally appro-

priate more often are less likely to watch the video, $\rho=-0.20$ ($p=0.06$). We find no significant correlation between the decision to watch the video and how often individuals consider it more morally appropriate to take the envelope or open the envelope first, $\rho=-0.01$ ($p=0.88$) and $\rho=-0.0691$ ($p=0.53$), respectively.

References

- [1] d’Adda, Giovanna, Yu Gao, Russell Golman and Massimo Tavoni (2018). “So Hot in Here: Information Avoidance, Moral Wiggle Room, and High Air Conditioning Usage.” Working paper.
- [2] Albrecht, Konstanze, Florentin Krämer and Nora Szech (2017). “Animal Welfare and Human Ethics: A Personality Study”. CESifo Working Paper No. 6609.
- [3] Allcott, Hunt and Judd B. Kessler (2019). “The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons.” *American Economic Journal: Applied Economics* 11 (1), 236–276.
- [4] Andreoni, James, and Lise Vesterlund (2001). “Which is the fair sex? Gender differences in altruism.” *Quarterly Journal of Economics* 116 (1), 293–312.
- [5] Andreoni, James, Justin M. Rao, and Hannah Trachtman (2017). “Avoiding the ask: a field experiment on altruism, empathy, and charitable giving”. *Journal of Political Economy*, 125 (3): 625–653.
- [6] American Society for the Prevention of Cruelty to Animals (2016). “New Research Finds Vast Majority of Americans Concerned about Farm Animal Welfare, Confused by Food Labels and Willing to Pay More for Better Treatment,” Press Release.
- [7] Bandura, Albert (2016). *Moral Disengagement: How People Do Harm and Live With Themselves*. New York: Worth Publishers.
- [8] Bartling, Björn, Florian Engl and Roberto A. Weber (2014). “Does Willful Ignorance Deflect Punishment? An Experimental Study.” *European Economic Review* 70, 512–524.
- [9] Bénabou, Roland, and Jean Tirole (2011). “Identity, Morals and Taboos: Beliefs as Assets.” *Quarterly Journal of Economics* 126(2): 805–55.
- [10] Bénabou, Roland, Armin Falk and Jean Tirole (2018). “Narratives, imperatives and moral reasoning.” NBER Working Paper 24798.
- [11] Bicchieri, Cristina and Erte Xiao (2009). “Do the Right Thing: But Only if Others Do So.” *Journal of Behavioral Decision Making* 22 (2), 191–208.

- [12] Broberg, Tomas, Ellingsen, Tore, and Magnus Johannesson (2007). “Is generosity involuntary?” *Economic Letters* 94 (1) 32–37.
- [13] Bundesministerium für Ernährung und Landwirtschaft [German Ministry of Nutrition and Agriculture] (2016). “Deutscher Ernährungsreport [German nutrition report]”
- [14] Cain, Daylain M. and Jason Dana (2012). “Paying People to Look at the Consequences of Their Actions.” Working paper.
- [15] Christie, R., and F. L. Geis. (1970) *Studies in Machiavellianism*. New York and London: Academic Press.
- [16] Cialdini, Robert B. and Noah J. Goldstein (2004). “Social Influence: Compliance and Conformity.” *Annual Review of Psychology* 55, 591–621.
- [17] Dana, Jason (2006). “Strategic Ignorance and Ethical Behavior in Organizations.” In *Ethics in Groups: Research on Managing Groups and Teams, Volume 8*, edited by Ann E. Tenbrunsel, 39-57. Bingley: Emerald Group Publishing.
- [18] Dana, Jason, Cain, Daylain M. and Robyn M. Dawes (2006). “What you don’t know won’t hurt me: Costly (but quiet) exit in dictator games.” *Organizational Behavior and Human Decision Processes* 100 (2), 193–201.
- [19] Dana, Jason, Roberto A. Weber, and Jason Xi Kuang (2007). “Exploiting Moral Wriggle Room: Experiments Demonstrating an Illusory Preference for Fairness.” *Economic Theory* 33(1): 67–80.
- [20] DellaVigna, Stefano, John A. List, and Ulrike Malmendier (2012). “Testing for altruism and social pressure in charitable giving.” *Quarterly Journal of Economics*, 127(1), 1–56.
- [21] Engel, Jannis and Nora Szech (2018). “A Little Good is Good Enough: Ethical Consumption, Cheap Excuses, and Moral Self-Licensing.” Working paper.
- [22] Exley, Christine (2015). “Excusing Selfishness in Charitable Giving: The Role of Risk.” *Review of Economic Studies* 83, 587–628.
- [23] Exley, Christine (2017). “Using Charity Performance Metrics as an Excuse Not to Give.” *Management Science*, forthcoming.

- [24] Exley, Christine and Judd Kessler (2018). “Motivated Errors.” Working paper.
- [25] Exley, Christine and Ragan Petrie (2018). “The Impact of a Surprise Donation Ask.” *Journal of Public Economics* 158, 152–167.
- [26] Falk, Armin, Neuber and Nora Szech (2019). XXX
- [27] Feiler, Lauren (2014). “Testing Models of Information Avoidance with Binary Choice Dictator Games.” *Journal of Economic Psychology* 45: 253–67.
- [28] Frey, Bruno and Stephan Meier (2009). “Social Comparison and Pro-social Behavior: Testing ‘Conditional Cooperation’ in a Field Experiment.” *American Economic Review* 94 (5), 1717–1722.
- [29] Freddi, Eleonora (2017). “Do People Avoid Morally Relevant Information? Evidence from the Refugee Crisis.” Working paper.
- [30] Ganguly, Ananda R., and Joshua Tasoff (2017). “Fantasy and Dread: The Demand for Information and the Consumption Utility of the Future.” *Management Science* 63 (12), 4037–4060.
- [31] von Gaudecker, Hans-Martin, Arthur van Soest and Erik Wengström (2011). “Heterogeneity in Risky Choice Behavior in a Broad Population.” *American Economic Review* 101 (2), 664–694.
- [32] Gneezy, Uri, Silvia Saccardo, Marta Serra-Garcia and Roel van Veldhuizen (2017). “Bribing the Self.” Working paper.
- [33] Golman, Russell, David Hagmann, and George Loewenstein (2017). Information Avoidance. *Journal of Economic Literature* 55 (1), 96–135.
- [34] Grossman, Zachary (2014). “Strategic Ignorance and the Robustness of Social Preferences.” *Management Science* 60 (11), 2659–2665.
- [35] Grossman, Zachary, and Joël van der Weele (2017). “Self-Image and Willful Ignorance in Social Decisions.” *Journal of the European Economic Association*, forthcoming.
- [36] Guerrero, Alexander A. (2007). “Don’t Know, Don’t Kill: Moral Ignorance, Culpability, and Caution.” *Philosophical Studies* 136, 59–97.

- [37] Hallsworth, M., John List, Robert Metcalfe and Ivo Vlaev. (2017). “The behavioralist as tax collector: Using natural field experiments to enhance tax compliance.” *Journal of Public Economics* 114, 14–31.
- [38] Hertwig, Ralph and Christoph Engel (2016). “Homo Ignorans Deliberately Choosing Not to Know.” *Perspectives on Psychological Science* 11 (3), 359–372.
- [39] Huck, Steffen, Nora Szech and Lukas Wenner (2018). “More Effort with Less Pay: On Information Avoidance, Belief Design and Performance.” KIT Working Paper.
- [40] Jackall, Robert (1988). *Moral mazes: The world of corporate managers*. Oxford: Oxford University Press.
- [41] Karlsson Niklas, George Loewenstein, and Duane Seppi. 2009. “The Ostrich Effect: Selective Attention to Information.” *Journal of Risk and Uncertainty* 38 (2): 95–115.
- [42] Krupka, Erin L. and Roberto A. Weber (2013). “Identifying Social Norms Using Coordination Games. Why Does Dictator Game Sharing Vary?” *Journal of the European Economic Association* 11, 495–524.
- [43] Larson, Tara, and C. Monica Capra (2009). “Exploiting Moral Wiggle Room: Illusory Preference for Fairness? A Comment.” *Judgment and Decision Making* 4(6), 467–474.
- [44] Lazear, Edward P., Ulrike Malmendier, and Roberto A. Weber (2012). “Sorting in Experiments with Application to Social Preferences.” *American Economic Journal: Applied Economics* 4(1), 1–36.
- [45] Litman, L., Robinson, J., and Abberbock, T. (2016). “TurkPrime. com: A versatile crowdsourcing data acquisition platform for the behavioral sciences.” *Behavior Research Methods*, 1–10.
- [46] Matthey, Astrid and Tobias Regner (2011). “Do I Really Want to Know? A Cognitive Dissonance-Based Explanation of Other-Regarding Behavior.” *Games*, 2, 114–135.

- [47] Miller, Suzanne M. (1987). “Monitoring and Blunting: Validation of a questionnaire to assess styles of information seeking under threat.” *Journal of Personality and Social Psychology*, 52(2), 345–353.
- [48] OECD (2018). *Behavioural Insights for Public Integrity Harnessing the Human Factor to Counter Corruption*. OECD Public Governance Reviews, OECD Publishing, Paris.
- [49] Onwezen, Marleen C. and Cor N. van der Weele (2016). “When indifference is ambivalence: Strategic ignorance about meat consumption.” *Food Quality and Preference* 52, 96-105.
- [50] Oster, Emily, Ira Shoulson, and E. Ray Dorsey (2013). “Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease.” *American Economic Review* 103 (2): 804–30.
- [51] Rabin, Matthew (1995). “Moral Preferences, Moral Constraints and Self-Serving Biases.” Working paper.
- [52] Rayner, Steve (2012). “Uncomfortable Knowledge: The Social Construction of Ignorance in Science and Environmental Policy Discourses” *Economy and Society* 41 (1), 107–25.
- [53] Rosen, G. (2003). “Culpability and ignorance.” *Proceedings of the Aristotelian Society*, 103 (1), 61–84.
- [54] Roth, A. (2007). “Repugnance as a Constraint on Markets,” *Journal of Economic Perspectives* 21 (3), 37–58.
- [55] Schultz, P. Wesley, Jessica M. Nolan, Robert B. Cialdini, Noah J. Goldstein and Vidas Griskevicius (2007). “The Constructive, Destructive and Reconstructive Power of Social Norms.” *Psychological Science* 18 (5), 429–434.
- [56] Shang, Jen and Rachel Croson. (2009). “A Field Experiment in Charitable Contribution: The Impact of Social Information on the Voluntary Provision of Public Goods.” *Economic Journal* 119, 1422–1439.
- [57] Sicherman, Nachum, George Loewenstein, Duane J. Seppi, and Stephen P. Utkus (2016). “Financial Attention.” *Review of Financial Studies* 29 (4): 863–897.

- [58] te Velde, Hein, Noelle Aarts and Cees van Woerkum (2002). “Dealing with Ambivalence: Farmers’ and Consumers’ Perceptions of Animal Welfare in Livestock Breeding.” *Journal of Agricultural and Environmental Ethics* 15, 203–219.
- [59] van der Weele, Joël (2014) “Inconvenient Truths: Determinants of Strategic Ignorance in Moral Dilemmas.” Working paper.
- [60] Zane, Daniel M., Julie R. Irwin and Rebecca Walker Reczek (2016). “Do less ethical consumers denigrate more ethical consumers? The effect of willful ignorance on judgments of others. *Journal of Consumer Psychology* 26 (3), 337–349.
- [61] Zimmerman, Michael J. (1997). “Moral responsibility and ignorance.” *Ethics*, 107 (3), 410–426.