

Inconvenient Truths:

A Note on Information Avoidance and the Price of Fairness*

Joël van der Weele (✉) Cristina Figueroa-Sisniega

October 16, 2025

Abstract

Previous literature has shown that people are often reluctant to learn whether individually profitable actions have negative consequences for others. In an experimental allocation decision, we vary the ‘inconvenience’ of becoming informed about the payoffs of another player by changing the costs and benefits of choosing the fair outcome. Making the fair allocation cheaper to implement turns out to have a multiplier effect, raising both altruistic choices of informed subjects and the fraction of subjects that chooses to become informed. Thus, in situations of uncertainty, subsidizing altruistic choices to decision makers could be an effective tool for raising social welfare. By contrast, variations in the size of recipients’ potential payoffs have a smaller effect on ignorance and fair choices.

JEL-codes: D83, C72, C91.

Keywords: willful ignorance, prosocial behavior, dictator games, experiment.

*Joël J. Van der Weele: University of Amsterdam, Tinbergen Institute. Email: j.j.vanderweele@uva.nl.
Cristina Figueroa-Sisniega: University of Amsterdam, Tinbergen Institute. Email: c.figueroa@uva.nl
Department of Economics, University of Amsterdam, Roeterstraat 11, 1018WB Amsterdam, the Netherlands.

“It is difficult to get a man to understand something when his income depends on his not understanding it.”

Upton Sinclair

1 Introduction

People appear to be generally reluctant to acquire information that may lead to difficult tradeoffs. For example, Ehrich and Irwin (2005) show that many consumers fail to inquire into ethical issues associated with cheap products, even though they will use the information when it is given to them. Norgaard (2006) concludes that subjects in her Norwegian sample avoid information about the reality of climate change in order to preserve established living patterns. In the laboratory, Dana *et al.* (2007, DWK hereafter) study an experimental situation in which a ‘dictator’ chooses to find out free of charge whether an action that carries a small cost to the dictator will help another person in the lab. Many dictators choose to ignore this information, and simply choose the action that gives them the most money, even though most dictators choose the fair outcome when they are (exogenously) informed about the benefits.

An important question is how and when the choice environment facilitates such information avoidance and the resulting detrimental outcomes for social welfare. Specifically, while the costs and benefits of altruism are known to be key determinants of altruistic decisions under certainty (Andreoni and Miller, 2002), it is not clear whether this is the case under uncertainty. Moreover, in an environment where subjects have the option to remain uninformed, the costs and benefits of different actions may affect both the decision to become informed and the subsequent allocation decision.

In this paper, we investigate how the costs and benefits of altruism or fairness affect the choice to learn the outcomes for others and the resulting altruistic choices. Using an experimental design based on DWK, we vary the cost to the dictator of implementing the fair allocation, making information more or less ‘inconvenient’. In addition, we vary the benefit to others that may result from a prosocial choice. These manipulations allow us to investigate the effect of the ‘price of fairness’ on information acquisition and prosocial decisions. To do so, we use both a laboratory study, as well as an online replication study.

The results show that under uncertainty, variations in the cost of fairness have a multiplier effect. First, lower costs cause informed subjects to behave more prosocially. Second, they substantially decrease the percentage of subjects who choose ignorance. These effects translate into a substantial increase in recipients’ payoffs and in efficiency. Furthermore, information management by dictators is asymmetric, as increasing the potential losses or gains of the passive recipient has little effect on ignorance or prosocial behavior, and results in more unequal allocations.

These findings show how ignorance is driven at least partially by a forward-looking process, which

is affected by monetary incentives of downstream decisions. One implication is that changing relative prices of actions or products may have large effects on behavior when there is uncertainty about the impact of these actions. In particular, our data corroborate the theory by Hestermann *et al.* (2020), that willful ignorance or self-deception generally should increase the price elasticity of ethically problematic products: higher prices decrease buying both through a standard substitution effect and through an increased willingness attend to the negative impact of the product.

This paper contributes to a fast-growing literature on information avoidance or willful ignorance (Sweeny *et al.*, 2010; Golman *et al.*, 2017; Hertwig and Engel, 2016). In particular, the current experiment is based on an experimental game developed DWK, which has produced numerous replications and follow-ups, summarized most recently in Vu *et al.* (2023). Several papers have looked into the mechanisms behind ignorance and its drivers (Trimmer *et al.*, 2021), for instance due to self-image (e.g. Grossman and Van Der Weele, 2017), social norms (Spiekermann and Weiss, 2016), complexity (Exley and Kessler, 2023) or the price of information (Serra-Garcia and Szech, 2022).

In the most closely related study, Feiler (2014) investigates ignorance while varying payoffs of the underlying game. She finds that small manipulations of the payoff structure do not lead to (significant) changes in behavior. These conclusions are based on a within subject design in which participants faced many (20) different games in succession, and incentives for each individual game were weak. By contrast, our study uses a between subject design in which each dictator made decision in only one, well-incentivized environment, using a maximally simple experimental setting that strips away all the noise. In a meta-study of several experiments, Vu *et al.* (2023) show no average effect of higher incentives for selfishness (i.e., the decision maker’s payoff advantage) on willful ignorance, but this conclusion depends on comparing payoffs across different experiments, which differ in the design and sample. Two recent papers have studied the effect of changes to product prices on willful ignorance, focusing on settings where products may have negative externalities, but do not find evidence that prices affect ignorance (Pace *et al.*, 2025; Amasino *et al.*, 2025). Thus, we add to this literature by providing the first evidence that ignorance responds to changes in payoffs of a subsequent decision.

We also contribute to a discussion about the motivations behind willful ignorance (Exley and Kessler, 2023; Vu *et al.*, 2023). While the multiplier effect of payoffs is consistent with standard models of social preferences, these models do not explain all aspects of the data. Evidence from a post-decision questionnaire shows evidence that both monetary payoffs and self-image concerns are important drivers of willful ignorance.

2 Experimental Design

The experimental setup is a generalization of the “hidden information treatment” in DWK. In all treatments subjects are randomly paired in groups of two, consisting of a ‘dictator’ (Player X) and a ‘recipient’ (Player Y). Both players are paid according to the decision of the dictator. The recipient is passive and does not make any decision and payoffs are denoted in terms of experimental currency (EC).

In the *Baseline* treatment, the dictator is facing the following situation. She can choose between two actions A and B , resulting in a payoff for the dictator of 100 or 60 EC respectively. Both players were told that before the start of the experiment a computer randomly determined the payoffs of the recipient associated with the dictator’s actions. Table 1 replicates the presentation of the possible payoffs given to the subjects in the instructions (see also the screenshots in Appendix B). Players were told that Game 1 and Game 2 were equally likely to be chosen by the computer. In the remainder of this paper, we will refer to Game 1 as the “Conflicting Interests Game” (CIG), since the dictator faces a trade-off between his own and the recipients’ payoff, and to Game 2 as the “Aligned Interests Game” (AIG), since A implements the highest payoffs for *both* players.

Player X chooses	Player X receives	Player Y receives
A	100	10
B	60	60

(a) Game 1

Player X chooses	Player X receives	Player Y receives
A	100	60
B	60	10

(b) Game 2

Table 1: The experimental games: the Conflicting Interests Game (CIG) on the left and the Aligned Interests Game (AIG) on the right. Each game has been chosen with 50% probability.

The dictator makes two choices. First, she has to decide whether to find out which game is being played. She faces a screen with a payoff matrix that shows her own payoffs, but where Player Y’s payoffs are replaced by a question mark. The screen features two buttons saying “Reveal game” or “Don’t reveal”. If the dictator decides to reveal the game she moves to the next screen where the full matrix is shown, as well as two buttons for choosing A and B . If she decides not to reveal, the question marks remain on the next screen. Thus, ignorance is not a default option as in DKW, but has to be chosen actively.¹ Once the information decision is made, the subject proceeds to choose A or B . There are no costs attached to remaining ignorant or acquiring information. To mimic most real-world information decisions, and in keeping with the previous experiments on this topic, the decisions in the experiment are made anonymously.

¹Grossman (2014) finds that changing from the default of ignorance to an active choice roughly halves the ignorance rate.

Most importantly, the recipient does not learn the dictator’s decision to reveal, and the dictator knows this.

Treatments. The experiment features three additional experimental conditions. The *Cheap Fairness* treatment is designed to test whether people are more or less likely to acquire information when it is cheaper to act prosocially. This treatment is equivalent to the *Baseline* treatment, except that the payoff of Player X associated with action A is now 70 EC (instead of 100 EC), so that the dictator has to give up less to implement an equal distribution in the CIG.

The remaining two treatments are designed to test whether subjects are more or less likely to remain ignorant when the expected loss to the other party of a self-interested choice increases. In the *Increased Loss* treatment, we vary the potential effect of self-interested behavior on the recipient. Payoffs are equivalent to the *Baseline* treatment, except that the worst possible outcome for Player Y is -20 EC (instead of 10 EC), i.e. the recipient may lose part of her show-up fee. In the *Increased Loss* treatment, the worst possible outcome for the receiver was losing these 20 points.

In the *Likely Loss* treatment, payoffs are the same as in the *Baseline* treatment, but the CIG is now likely to be selected with a probability of 0.8 instead of 0.5. This probability is explicitly mentioned a few times throughout the instructions.

Procedures. We conducted two experiments using this design. The first was a laboratory experiment conducted in 2014 at the Frankfurt Laboratory for Experimental Economics (FLEX). Participants were recruited through ORSEE (Greiner, 2003) from the subject pool at Goethe University Frankfurt, which includes students from a wide range of academic disciplines. Each participant received a 4 euro show-up fee and completed the experiment at individual computer terminals. The EC-to-currency exchange rate was $10 \text{ EC} = 1 \text{ euro}$.

The second experiment was a pre-registered (AsPredicted#216527) replication of the first experiment, and it was conducted online in 2025 via Prolific. We recruited participants globally, provided they were fluent in English. The sample size for this experiment was determined by a power analysis, based on the laboratory results. Online participants received a 1.50 pound show-up fee and the EC-to-currency exchange rate was $80 \text{ EC} = 1 \text{ pound}$. Apart from the setting and the exchange rate, the Online experiment slightly modified the *Increased Loss* treatment because deducting money from the participation fee is not allowed in Prolific. Thus, in order to generate the increased loss effect (worst possible outcome for Player Y is -20 EC) we endowed every participant in the online experiment with 20 EC as a reward for answering a short questionnaire at the end of the experiment. The -20 payoff for Player Y in the *Increased Loss* meant that Player Y “lost” the 20 EC endowment. Screenshots of the instructions and experimental interface for both experiments are provided in Appendix B.

3 Results

We begin by analyzing the patterns observed in the pooled dataset, which combines the lab data ($n = 132$ dictators, collected in 2014) and the online data ($n = 242$ dictators, collected in 2025). The comparison between the two subsamples follows in the next subsection. Descriptive statistics for each sample are provided in Table A2 in Appendix A.

Pooled Data

We start by investigating the ‘ignorance rate’ – the rate at which participants choose not to know which game they are playing. Note that ignorance makes it impossible to choose the fair or equitable outcome intentionally: in the absence of information about the game being played, either option may benefit or hurt the recipient.² As shown in the left panel of Figure 1, the Pooled Data shows that 32% of the subjects remained ignorant in the *Baseline* treatment, as indicated by the light gray area. This ignorance rate drops significantly to 16% in the *Cheap Fairness* treatment ($Z = 2.628$, $p = 0.009$, 2-tailed Z-test (ZT)). There is also a decrease to 24% in both the *Increased Loss* ($Z = 1.315$, $p = 0.189$, ZT) and the *Likely Loss* treatments ($Z = 1.240$, $p = 0.215$, ZT), although these differences are not statistically significant.

A more comprehensive measure of the pursuit of self-interest sums the ignorance rate and the fraction of subjects who chose option *A* in the CIG, after learning that this decision negatively affected another participant (the blue areas in Figure 1). In the *Pooled Data*, the share behaving inconsistently with fairness in the *Baseline* treatment is 58%. This proportion significantly decreases to 35% in the *Cheap Fairness* treatment ($Z = 3.208$, $p = 0.001$, ZT). The levels in the *Increased Loss* (52%) and *Likely Loss* (61%) treatments are not statistically different from the Baseline ($Z = 0.908$, $p = 0.364$, ZT; $Z = 0.383$, $p = 0.702$, ZT; respectively).

²Thus, choosing option B after choosing ignorance is dominated action, and choosing it seems to reflect a misunderstanding of the game. Indeed, a choice of ignorance is followed by a choice of option B in only 9% of cases. Nevertheless, in Appendix B, we provide a modified version of Figure 1 that excludes such dominated choices. Instead of overall ignorance rates, it shows the proportion of participants who remained ignorant and chose option A. The observed patterns are very similar to those in the main text, reinforcing our interpretation that ignorance is predominantly associated with selfish choices.




Figure 1: Ignorance and self-interested behavior by treatment. “Remained ignorant” is the fraction of subjects choosing “Don’t Reveal”. “Chose A in CIG” is the fraction of subjects who revealed, faced the CIG game and chose A over the total number of subjects in the treatment. The total height of the bar thus measures all subjects whose behavior is not compatible with a desire to implement a fair outcome. The total number of observations in each treatment is at the bottom of the bar.

Table 2 summarizes dictators’ choices and the resulting payoffs for the pooled data (Table A3 in Appendix A presents the corresponding table for each sample, separately). Comparing the *Baseline* with the *Cheap Fairness* treatment in Table 2 shows that an increase in the cost of prosocial behavior delivers a ‘double whammy’: it doubles the ignorance rate *and* it increases self-interested behavior by informed dictators in the CIG by approximately 65% (A choices by dictators who were informed that they played the CIG in the *Baseline* are 64.1% versus 39.1% in the *Cheap Fairness* treatment $Z = 2.295$, $p = 0.022$, ZT). As a result of this multiplier effect, the payoffs of the recipient in the *Baseline* treatment are only about half those in the *Cheap Fairness* treatment (56% to be precise), and overall altruism, defined as the percentage of B choices among the dictators in the CIG, drops by approx. 40% in the *Baseline* relative to the *Cheap Fairness* condition ($Z = 3.675$, $p = 0.000$, ZT).

	<i>Baseline</i>	<i>Cheap Fairness</i>	<i>Increased Loss</i>	<i>Likely Loss</i>	Informed subjects	Ignorant subjects
Ignorance (CIG + AIG)	32.3%	16.0%	23.6%	24.2%	-	100.0%
A choices (CIG informed)	64.1%	39.1%	67.6%	61.4%	57.5%	-
Avg. Dictator Payoff CIG (points)	90.8	64.3	89.2	87.2	80.0	95.0
Avg. Recipient Payoff CIG (points)	21.5	38.4	1.7	26.0	27.0	8.2
Overall Altruism (B choices CIG)	23.0%	56.9%	27.1%	32.0%	42.5%	7.1%

Table 2: Behavior and average payoffs in points. The first row shows the choices of ignorance as a proportion of all choices. The second row reports the percentage of subjects who chose option *A* among those who revealed and learned they were in the CIG. The middle two rows show average payoffs (net of the show-up fee) of those who played the CIG, either knowingly or not. The final row shows the number of *B* choices as a percentage of all choices in the CIG. The last two columns include data from all treatments.

Second, while the ignorance rate in the *Increased Loss* and the *Likely Loss* treatments decreases relative to the *Baseline* (although insignificantly as noted before), the rate of choosing *A* among informed dictators in the CIG does not decrease. In the *Increased Loss*, *A* choices for subjects who are informed that they are in the CIG increases somewhat relative to the baseline, while there is slight decrease in the case of the *Likely Loss* treatment, but none of these differences are statistically significant ($Z = 0.318$, $p = 0.750$, ZT and $Z = 0.268$, $p = 0.788$, ZT, respectively), nor do these treatments generate higher levels of overall altruism (fraction of *B* choices in the CIG) relative to the *Baseline* treatment ($Z = 0.496$, $p = 0.620$ for the first, and $Z = 1.170$, $p = 0.242$ for the second ZT comparisons, respectively). Thus, changing recipient payoffs does not significantly alter behavior relative to the *Baseline*. All our results are robust in regressions, which include a dummy for the origin of the data (laboratory or online), as well as gender – see Tables A7 - A9 in Appendix A.

Finally, when we compare the behavior of informed and uninformed dictators in all treatments (the final two columns of Table 2), we see that approximately 42% (76 out of 179) of informed dictators choose action *B* in the CIG. By contrast, only 7% (4 out of 56) of the subjects who remained ignorant chose *B*.³ Inequality therefore increases with ignorance: whereas informed dictators earn about 2.9 times as much as the recipient, ignorant dictators earn about 11.6 times as much. Moreover, the ignorant dictator earns 18.75% more on average than the informed dictator.

In summary, we find that compared to the *Baseline*, the *Cheap Fairness* treatment results in lower ignorance and fairer outcomes. While the Loss treatments result in slightly (but not

³In addition, all but one of those who informed themselves and played the AIG choose *A* in the full sample. This indicates that there was a low amount of ‘noise’ in the experiment and that subjects understood the game well.

statistically significantly) lower ignorance rates, they have little impact on overall prosociality.

Laboratory vs. Online data

In this section we focus on the treatment patterns in the laboratory experiment (“Lab”) and the Online experiment (“Online”), displayed in Figure 1 (the second and third panels). Results are not statistically significant at conventional levels on one-sided tests, unless otherwise reported. Tables A4 - A6 in the Appendix provide the relevant test statistics.

We begin by comparing the *Baseline* and *Cheap Fairness* conditions. In both the Lab and Online samples, lowering the cost of fairness leads to a reduction in ignorance, though the effect is stronger in the Lab. There, ignorance drops from 31% to 6% (with statistical significant at $p = 0.011$), while in the Online sample it declines from 33% to 21%, though this difference is not statistically significant ($p = 0.123$). We observe a similar pattern in the overall share of behavior consistent with selfishness, which falls significantly in the *Cheap Fairness* condition compared to *Baseline* in both samples (Lab: 52.78% to 21.88%; Online: 61.67% to 41.94%). The same story holds for informed dictators in the CIG: the proportion choosing the selfish option (A) falls significantly, from 61.54% to 31.25% in the Lab sample and from 65.38% to 43.33% in the Online sample.

These results illustrate the multiplier effect in both samples (lowering the cost of altruism reduces both ignorance and selfish choices) but the effect seems to be more pronounced in the Lab. One likely explanation for the latter lies in the monetary incentives: the cost of being fair is relatively lower in the lab due to a more favorable exchange rate (20 EC = €1) than in the online sample (80 EC = £1). Additionally, differences in attention and engagement between lab and online environments may contribute to the differences in treatment effects.

Turning to the *Increased Loss* condition, we find that both samples exhibit similar, but statistically insignificant, reductions in ignorance compared to the *Baseline* (Lab: 30.56% vs. 16.67%; Online: 33.33% vs. 27.12%). Regarding the proportion of selfishness consistent behaviour, we find that for both samples the *Baseline* and the *Increased Loss* treatments show no significant differences (Lab: 52.78 % and 46.67%, respectively; Online: 61.67% and 54.24%). Likewise, the proportion of selfish choices by informed CIG dictators remains stable, with no significant differences (Lab: 61.54% vs. 64.29%; Online: 65.38% vs. 69.57%).

In the *Likely Loss* condition, ignorance in the Lab sample remains virtually unchanged compared to *Baseline* (29.41% vs. 30.56%), while in the Online sample it declines, though not significantly (33.33% vs. 21.31%). Regarding the proportion of behavior that is consistent with selfishness between *Baseline* and *Likely Loss*, these proportions remain comparable across treatments for both samples (Lab: (52.78% vs. 55.88%, ; Online: 61.67% in *Baseline* vs. 63.93% in *Likely Loss*). Among informed CIG dictators, selfish choices are statistically unchanged relative to the *Baseline* in both samples (Lab: 61.54% vs. 50.00%; Online: 65.38% vs. 66.67%).

Comparing the *Cheap Fairness* and *Loss* treatments, we cannot always reject the null hy-

pothesis for equality of the ignorance rate (see Tables A4 - A6 in the Appendix). However, unlike in the *Cheap Fairness* condition, where the overall share of behavior consistent with selfishness decreases relative to the *Baseline*, no such reduction is observed in the *Loss* conditions. Indeed, selfish behavior in the *Loss* treatments is higher than in the *Cheap Fairness* treatment, with statistical significance in the lab and the pooled samples. In other words, ignorance appears to have a multiplier effect only in the *Cheap Fairness* condition.

Since demographics across the Lab and Online samples may vary, we report treatment effects from regressions that control for gender (the demographic variable available in both samples) in Tables A7 - A9 in Appendix A. The results discussed in the main text are robust to including this control. We also find that males are more prone towards ignorance and selfishness than females.

Discussion

Andreoni and Miller (2002) have shown that altruistic actions respond in predictable ways to the price of altruism. The results in this paper show that the price of altruism also affects whether people want to learn the consequence of altruistic or prosocial actions. Subjects in the experiment are more motivated to look the other way if prosocial or altruistic behavior is costly, consistent with a desire to avoid more ‘inconvenient’ facts.

The fact that ignorance responds to incentives in downstream decisions, also indicates that it stems from a ‘rational’, or at least forward-looking decision making process. Below, we discuss which motives may be at play in such a process, as well as the implications of our results.

Underlying motives. Well-known models of social preferences like those of Fehr and Schmidt (1999) posit that subjects trade-off their own payoffs and the wish to implement a fair or equal outcome. These models can capture some of our results. Subjects motivated to avoid inequality will choose information and choose B in the CIG. Those who do not care (enough) about fairness will choose A in both the CIG and the AIG, and are indifferent between information and ignorance. When the fair option becomes cheaper, it becomes an attractive option to more subjects, which can explain why more subjects become interested in information. However, if subjects trade-off fairness with their own payoffs, it is hard to explain why they barely react to the payoffs of the other player.

An alternative explanation is that ignorance serves as an ‘excuse’ for selfish behavior, either towards others or towards oneself. For instance, Grossman and Van Der Weele (2017) argue that willfully chosen ignorance is a compromise between material interests and a desire to maintain a (self-)image as an altruistic or fair-minded individual. While this model can explain findings that social preference models cannot capture (like a willingness to pay for ignorance, and an increase in selfish behavior compared to forced information treatments), it gives similar predictions to

social preference models in the setup of this paper.⁴

To get a better understanding of participants’ motives, we analyze responses to a post-experimental questionnaire. In the online experiment, dictators were asked to indicate (through a single-choice item question) the main reason behind their decision to reveal or not reveal the payoffs. Those who chose to reveal could select from four options: (1) concerns about payoffs (consistent with an outcome-based motive), (2) concerns about how they would feel about themselves (consistent with an image-based motive), (3) curiosity, or (4) none of the above. For those who chose not to reveal, the same options were available except for curiosity. In the Lab experiment, participants answered an open-ended question: “Did you reveal the payoffs of Player Y? Why (not)?” We used a large language model (OpenAi’s GPT model 4o) to classify these open-ended responses into the same motivation categories used in the online experiment. The distribution of motivations across both experiments, by reveal decision, is presented in Figure 2 below.

⁴Explicit proofs to this extent appear in Van der Weele (2012). In particular, the theory predicts that the drop in ignorance in the *Cheap Fairness* treatment occurs because obtaining a good self-image is now cheaper, and information therefore less threatening. The model also predicts a drop in ignorance in the *Increased Loss* treatment, since dictators now feel more guilty when choosing self-interestedly.




Figure 2: Motivations for revealing or not revealing the payoffs, based on responses to a single-choice question (with multiple predefined options) in the online experiment and an open-ended question in the lab experiment. Responses were categorized into four motive types: image-based, outcome-based, curiosity, and other. Note that the “I was curious” option was only available to dictators who chose to reveal in the online experiment. The open-ended responses from the lab were classified into the same categories and then merged with the online data, as explained in the text.

The motivations behind choosing ignorance thus reflect a concern for both money and (self)-image. Supporting the image-based model, 34.4% of dictators who chose not to reveal cited self-image concerns as their primary reason. At the same time, 30% of those who remained ignorant reported being mainly concerned about the actual outcomes and payoffs, consistent with the outcome-based model, where subjects care predominantly about their own payoff.

In contrast, revealing behavior appears to be more strongly driven by outcome-based motivations: among those who revealed, the most commonly cited reason was concern about payoffs (46.1%). Curiosity was also a significant factor, with 24.6% citing it as their main motive—suggesting that some individuals may have an intrinsic preference for information.

Self-image concerns are cited by 18% of subjects.

Interestingly, selection of the “None of the above” option was substantially higher among those who chose not to reveal the payoffs: 35.6% compared to just 11.3% among those who did reveal. This may indicate that motivations beyond image and outcome concerns play a role in the decision to remain ignorant. However, it is also possible that choosing this option serves as a form of excuse, i.e. they want to obfuscate their choice to the experimenter. This interpretation aligns with previous findings in Saccardo and Serra-Garcia (2023) who show that individuals who seek cognitive flexibility in order to enable self-serving behavior often provide vague explanations for their choices that fall under “other reasons”, a deliberately non-transparent category.

Finally, we also examined the distribution of motives by reveal decision across treatments, as shown in Figure A1 in the Appendix. According to a Chi-square test of independence, the distribution of motives among those who chose not to reveal does not differ significantly across treatments ($\chi^2 = 9.27, df = 6, p = 0.159$). In contrast, the distribution of motives among those who did reveal does differ significantly by treatment ($\chi^2 = 19.89, df = 9, p = 0.019$). Figure A1 suggests that the distribution of motives for revealing in the Cheap Fairness treatment is somewhat distinct: individuals appear to be less driven by image concerns and more motivated by payoff considerations compared to other treatments. However, this observation should be treated with caution due to the small sample size.

Implications. The current results concern an abstract decision in a highly stylized setting, and a very stark information choice. In many consumption decisions, the choice for information or ignorance is likely to be less explicit than in the experiment. This makes it harder to argue that ‘one could have known’, and easier to use ignorance as an excuse, either to oneself or to others. As a consequence, we conjecture that strategic ignorance is at least as prevalent as in the lab. Indeed, a handful of studies have conducted tests of willful ignorance in applied settings, and found clear evidence in favor (Ehrich and Irwin, 2005; Epperson and Gerster, 2024; Woolley and Risen, 2021), though other studies have found mixed evidence (Pace *et al.*, 2025; Amasino *et al.*, 2025).

Assuming some external validity, our results have implications for corporate governance and policy making in contexts with payoff uncertainty. First, the results of the *Increased Loss* and *Likely Loss* treatments suggest that potentially large negative consequences for others do little to encourage information acquisition. This may help explain why large-scale frauds in organizations like Enron and state-sponsored atrocities carried out with the involvement of many may nevertheless remain ‘secret’ for a long time (Cohen, 2001). In the corporate sphere, such secrets may be discouraged by making managers directly responsible for the quality of the information they have (and report) about the organization. In fact, one of the provisions in the Enron-inspired Sarbanes-Oxley Act places the responsibility for the accuracy of information in financial reports with the executives.

Second, the results of the *Cheap Fairness* treatment show that in the presence of payoff

uncertainty, incentive schemes that reward prosocial actions will have a multiplier effect: not only do they influence behavior of informed agents, they also make individuals more perceptive towards opportunities for doing good. In the consumer domain, this is an argument for subsidies that lower the price of ‘ethical’ (e.g. fair-trade or eco-labeled) goods as it makes people more willing to learn about the social benefits of such products. Conversely, subsidies for socially or environmentally damaging products such as meat may make people less willing to learn about vegetarian alternatives (Hestermann *et al.*, 2020). Similarly, as pointed out by Upton Sinclair in the opening quote, organizational incentive schemes that reward on the basis of narrow performance criteria may induce strategic ignorance towards ethical violations that distract from these criteria.

Acknowledgements

We would like to thank Ferdinand von Siemens, Samuel Bowles, Roberto Weber, Zachary Grossman and Leonie Gerhards for comments, Zachary Grossman again for generously sharing his z-Tree code with me, and Sebastian Schäfer, Matthias Heinz, Karin Hettwer and Julija Kulisa for helping out with practical matters. We are indebted to the University of Amsterdam and the Vereinigung der Freunde und Förderer der Goethe-Universität for financial support. All errors are ours.

References

- AMASINO, D. R., OOSTERWIJK, S., SULLIVAN, N. J. and VAN DER WEELE, J. (2025). Seeking or ignoring ethical certifications in consumer choice. *Ecological Economics*, **229**, 108467.
- ANDREONI, J. and MILLER, J. (2002). Giving according to garp: An experimental test of the consistency of preferences for altruism. *Econometrica*, **70** (2), 737–753.
- COHEN, S. (2001). *States of Denial*. Cambridge: Blackwell.
- DANA, J., WEBER, R. R. and KUANG, J. X. J. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, **33** (1), 67–80.
- EHRICH, K. R. and IRWIN, J. R. J. (2005). Willful ignorance in the request for product attribute information. *Journal of Marketing Research*, **XLii** (August), 266–277.
- EPPELSON, R. and GERSTER, A. (2024). Willful ignorance and moral behavior. *Available at SSRN 3938994*.
- EXLEY, C. L. and KESSLER, J. B. (2023). Information avoidance and image concerns. *The Economic Journal*, **133** (656), 3153–3168.
- FEHR, E. and SCHMIDT, K. (1999). A Theory of Fairness, Competition and Cooperation. *The Quarterly Journal of Economics*, **114** (3), 164–817.
- FEILER, L. (2014). Testing models of information avoidance with binary choice dictator games. *Journal of Economic Psychology*, **In Press**.
- GOLMAN, R., HAGMANN, D. and LOEWENSTEIN, G. (2017). Information avoidance. *Journal of economic literature*, **55** (1), 96–135.
- GREINER, B. (2003). An Online Recruitment System for Economic Experiments. *Forschung und wissenschaftliches Rechnen*, **63**, 79–93.
- GROSSMAN, Z. (2014). Strategic ignorance and the robustness of social preferences. *Management Science*, **60** (11), 2659–2665.
- and VAN DER WEELE, J. J. (2017). Self-image and willful ignorance in social decisions. *Journal of the European Economic Association*, **15** (1), 173–217.
- HERTWIG, R. and ENGEL, C. (2016). Homo ignorans: Deliberately choosing not to know. *Perspectives on Psychological Science*, **11** (3), 359–372.
- HESTERMANN, N., LE YAOUANQ, Y. and TREICH, N. (2020). An economic model of the meat paradox. *European Economic Review*, **129**, 103569.
- NORGAARD, K. (2006). ”We don’t really want to know”. Environmental Justice and Socially Organized Denial of Global Warming in Norway. *Organization & Environment*, **19** (3), 347–370.
- PACE, D. D., IMAI, T., SCHWARDMANN, P. and VAN DER WEELE, J. J. (2025). Uncertainty about carbon impact and the willingness to avoid co2 emissions. *Ecological Economics*, **227**, 108401.
- SACCARDO, S. and SERRA-GARCIA, M. (2023). Enabling or limiting cognitive flexibility? evidence of demand for moral commitment. *American Economic Review*, **113** (2), 396–429.
- SERRA-GARCIA, M. and SZECH, N. (2022). The (in) elasticity of moral ignorance. *Management Science*, **68** (7), 4815–4834.
- SPIEKERMANN, K. and WEISS, A. (2016). Objective and subjective compliance: A norm-based explanation of ‘moral wiggle room’. *Games and Economic Behavior*, **96**, 170–183.
- SWEENEY, K., MELNIK, D., MILLER, W. and SHEPPERD, J. (2010). Information avoidance: Who, what, when and why. *Review of General Psychology*, **14** (4), 340–353.
- TRIMMER, P. C., McELREATH, R., AUSTER, S., BROWN, G. D., DANA, J., GIGERENZER, G., GOLMAN, R., HILBE, C., KANDLER, A., KAREEV, Y. *et al.* (2021). The zoo of models of deliberate ignorance.
- VAN DER WEELE, J. J. (2012). When Ignorance is Innocence: On Information Avoidance in Moral Dilemmas. *SSRN working paper*, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1844702.
- VU, L., SORAPERRA, I., LEIB, M., VAN DER WEELE, J. and SHALVI, S. (2023). Ignorance by choice: A meta-analytic review of the underlying motives of willful ignorance and its consequences. *Psychological Bulletin*, **149** (9-10), 611.
- WOOLLEY, K. and RISEN, J. L. (2021). Hiding from the truth: When and how cover enables information avoidance. *Journal of Consumer Research*, **47** (5), 675–697.