The lesser of two evils: Explaining a bad choice by revealing the choice set

Andras Molnar^{1*} & Shereen J. Chaudhry¹

¹Department of Social and Decision Sciences, Carnegie Mellon University

February 13, 2020

WORD COUNT: 9,296

*Corresponding author. Email: andrasm@andrew.cmu.edu; Tel.: (412) 478 7046; Address: 5000 Forbes Av., 321 Porter hall, Pittsburgh, PA 15213

Author contributions & acknowledgements: AM and SJC designed the experiments. AM collected and analyzed the data. AM and SJC wrote the manuscript. The authors wish to thank George Loewenstein and Alex Van Zant for their valuable suggestions for the experimental design and their feedback on earlier versions of the manuscript, Russell Golman, Adam Greenberg, Howard Kunreuther, Daniel Oppenheimer, and Silvia Saccardo for their comments on earlier versions of the manuscript.

Funding & ethical review: The study was funded by Carnegie Mellon University. The study was approved by the Institutional Review Board of Carnegie Mellon University (HS15-493).

Author's note: Molnar is PhD student in the Social and Decision Sciences department at Carnegie Mellon University, Email: andrasm@andrew.cmu.edu. Chaudhry is currently affiliated with the Booth School of Business, University of Chicago, Email: shereen.chaudhry@chicagobooth.edu.

THE LESSER OF TWO EVILS

Abstract. Making the right choice sometimes involves selecting the “lesser of two evils,” and only seeing the chosen option can lead others to misunderstand the decision maker’s intentions. Are decision-makers intrinsically driven to fix this misjudgment by revealing the choice set? If so, why, and what is the effect on the audience? Previous studies could not examine this desire to be understood because the research designs did not isolate the decision to reveal information from the original choice. In two experiments (N=448 pairs), we show that people are willing to pay *ex post* to reveal their choice set to the recipient, even after a one-shot anonymous interaction with no reputational consequences, and in some cases even when doing so reveals their selfish intentions. We find that this revealing behavior is effective at improving recipients’ rating of their outcome when it signals generous intentions, but not when it signals selfish intentions. The choice to reveal is driven by concern for the beliefs of strangers, but *only* when revealing signals generous intentions; those who reveal a choice that appears selfish report doing so out of a desire to be or appear honest. And though some people leave a misunderstanding in place when it is self-enhancing to do so, almost no one is willing to *create* a misunderstanding (by hiding the other option), even when it could conceal selfish behavior. **(226 words)**

JEL classification C91, D64, D82, D83, D84

Keywords Beliefs, Correspondence bias, Impression management, Information disclosure, Social image

INTRODUCTION

People often find themselves in situations where they must choose from a constrained set of options, each of which, when evaluated separately, seems like a terrible choice. In such cases, despite their best intentions, decision-makers are perceived to have done the wrong thing, especially when observers have incorrect beliefs about the set of available options. In addition to famous social and moral dilemmas (e.g., the trolley problem, Sophie's choice, the overcrowded lifeboat), examples for such situations exist in almost all domains of everyday life: prioritizing which patients to save first, choosing between risky medical treatments, allocating budget cuts, making decisions about layoffs, rejecting outstanding faculty candidates when there is only one position, assigning boring and unrewarding tasks that someone nevertheless has to do, deciding which internship to accept when they are all terrible, choosing a meal late at night when most places are closed, and choosing between two massively unpopular candidates running for office. With regard to the last example, during the presidential election of 2016, pundits and voters often and explicitly described it as choosing "the lesser of two evils," since the unprecedented low approval rate of both candidates forced voters to make a very difficult choice (Fabrizio, 2016).

(Mis)attribution and the correspondence bias. These choices can be difficult and frustrating to make, even when the decision-maker, in absence of an audience, has a clear preference over options—why? One reason is that people do not want their actions to be misinterpreted by others; they do not want to be misunderstood or misjudged. Observers are eager and quick to make inferences about underlying dispositions and preferences from observed behavior and choice (Heider & Simmel, 1944; Jones & Davis, 1965; Nisbett & Ross, 1991), often overattributing to dispositional factors and underestimating situational influences (Quattrone, 1982). Context is often an important, if not the primary, determinant of behavior (Mischel, 1977), and as a result, incomplete information about the context of a choice can cause onlookers to make unwarranted judgments about a decision-maker's intention and character. Indeed, judging a person based on her behavior can lead to both mistaken beliefs about her preferences (Jones & Harris, 1967) and also mistaken predictions of her future behavior (Bierbrauer, 1979). An extreme version of this mistake is the correspondence bias when observers, completely neglecting the causal role of the situational constraints, incorrectly draw inferences about the decision-maker's underlying characteristics and dispositions from behavior that can be *fully*

THE LESSER OF TWO EVILS

explained by the particular context in which they occur (Gilbert & Malone, 1995). Many of the examples provided earlier fall into this category: In most cases, a manager does not lay employees off because she has a perverse preference for firing people, a doctor does not postpone a patient's care out of personal discrimination, and many people who voted for Trump or Clinton did so because they strongly wanted to avoid electing the other candidate, not because they actively supported the person they voted for. As Miller & Nelson (2002) pointed out, observers tend to interpret others' actions as approach-motivated (i.e., choosing the option they like the most), even when they know that their own choices are motivated by avoidance (i.e., choosing the option they dislike the least).

Avoiding misattribution by revealing the choice set. In such "lesser of two evils" situations, decision-makers can avoid misjudgment by others by disclosing their full choice set, conveying that they were constrained in some way (e.g., the company board required them to lay off 40% of employees; they had to prioritize among patients who were all in critical condition; their department could only afford to hire one new faculty member). In correcting the misjudgment, revealing could help maintain a desirable social image, assuage the negative feelings of others, or simply convey the truth of their intentions (good or bad). In this sense, revealing the choice set can be considered a special case of signaling preferences and intentions: Instead of sending signals directly through the choices they select, decision-makers in this case send a message that is not directly payoff relevant (i.e., does not affect the welfare of others) (Golman, 2016). Rather, such a message draws the observers' attention to the context or choice set, which then allows them to adjust their dispositional inferences about the decision-maker in light of the newly discovered situational factors (Trope & Liberman, 1993).

People may want to avoid misjudgment for strategic (extrinsic) or non-strategic (intrinsic) reasons. When people are expecting to interact with the observer in the future (providing a possibility of reward or punishment), or when their reputation is at stake, they have a *strategic* motive to boost their image by revealing their choice set. This also means that they would choose not to reveal their choice set, if it could make them look worse. On the other hand, if a person gets some intrinsic or psychological benefit from revealing the choice set, they may do this even in anonymous, one-shot interactions, when reputation should not matter. Furthermore, depending on the underlying intrinsic drive, they may choose to reveal even if it conveys selfishness or some other negative attribute.

THE LESSER OF TWO EVILS

Previous work that has examined choice-set revealing (or hiding) behavior has done so with the presumption of its being a tool primarily to obtain a positive (or prevent a negative) image in the eyes of others, especially with respect to perceptions of generosity and fairness (Cappelen, Halvorsen, Sørensen, & Tungodden, 2017; Dana, Cain, & Dawes, 2006; Dana, Weber, & Kuang, 2007; Grossman & van der Weele, 2017). In contrast, we leave open the possibility for other motivations, and we investigate what those motivations are. Furthermore, in most of those previous studies, revealing behavior was often tied to the original decision itself: Either it was known before the decision was made that there would be a chance to reveal the choice set, or after the option to reveal was unveiled, the decision maker could return to alter their initial decision. Effectively, most studies examined the *ex ante* choice to reveal. This is important because, when the decision-maker knows in advance (before making a choice) that she will be able to provide excuses, apologies, or reveal intentions *after* implementing her choice, she can (and will) change her behavior accordingly (Andreoni & Rao, 2011; Ho, 2012). What is common in these studies is that the decision-maker's behavior and the observer's initial beliefs upon learning about the decision-maker's behavior are interdependent, since at the time of decision the decision-maker has control not only over her choice but also over what the observer will know about that choice. Thus, fully disentangling the choice and decision to reveal the choice set is essential if we want to study the intrinsic motive to have choices understood, ruling out strategic, payoff relevant motives.

In contrast, we have people make a decision without knowing about the opportunity to reveal. Later, or *ex post*, we give participants the opportunity to reveal (without the ability to change their original decision). In this way, the choice to reveal is only made afterwards and cannot influence the original decision. To our knowledge, no previous research has examined such *non-strategic ex post* revealing behavior, when the choice has been already made, observed, and is beyond the decision-maker's control. Are people willing to incur a cost to reveal information in such cases? And if so, what is their main motivation to do so? Similar to previous work, we believe that the motivation to reveal requires caring about the beliefs of others. (In our design, the decision maker is fully aware of their choice and its material consequences, so we rule out the types of beliefs about the self and self-signaling that do not depend on the beliefs of others.) However, this includes a variety of beliefs and a variety of reasons for caring about those beliefs.

THE LESSER OF TWO EVILS

In the next section, we review several of these potential *intrinsic* motivations that might be behind the drive to reveal the choice set.

Impression management and social image. Perhaps the account most frequently cited as an explanation for why people prefer to manage what is known about their choice set is to maintain a good impression in the eyes of others. People tend to judge, and be judged by, others along two universal dimensions (warmth and competence, see Fiske, Cuddy, & Glick, 2007), so they engage in a multitude of behaviors trying to maintain a positive impression: help others, be polite, act generously, behave competently, follow rules and social norms. Often times, they even incur costs to do so. Such concerns for a good impression and social image are not limited to strategic situations where reputation matters (i.e., involve direct or indirect reciprocity), so it seems that people have an intrinsic desire to uphold a favorable social image of themselves, which has been documented in many laboratory studies featuring one-shot, anonymous interactions lacking the threat of punishment or the promise of reward (Andreoni, Rao, & Trachtman, 2017; Cappelen et al., 2017; Dana et al., 2006, 2007; Grossman & van der Weele, 2017).

Concern for the feelings or expectations of others. A second account that could explain the desire to reveal the choice set focuses on decision-maker's concern for the feelings of others, or beliefs-based altruism (te Velde, 2018). A person's beliefs could lead them to feel a variety of emotions, such as disappointment, surprise, anger, curiosity, satisfaction, or relief. Instead of being motivated by to instill a positive image of themselves in the eyes of others, the decision maker might be motivated to inspire beliefs that lead to positive, rather than negative, emotions in the other person. If a choice between "two evils" results in a bad outcome for another person, revealing the choice set might move the affected party from thinking that the world is a harsh place, resulting in sadness or disappointment, to thinking that others do their best to be generous, resulting in gratitude or pleasant surprise. A more self-serving explanation for preventing negative feelings in others, especially feelings of disappointment, could stem from wanting to alleviate one's own negative feelings (e.g., guilt). This is the central element in the theory of guilt aversion, which assumes that people incorporate others' beliefs (and even second-order beliefs) into their utility functions, and they suffer disutility when their actions fail to meet others' expectations (Battigalli & Dufwenberg, 2007). This could lead to behavior that expresses a preference for not letting others feel frustrated and disappointed (Charness & Dufwenberg, 2006; Dufwenberg & Gneezy, 2000). Regardless of the source of concern for others' feelings,

THE LESSER OF TWO EVILS

this remains an additional, alternative explanation for revealing a choice set in a “lesser of two evils” scenario.

Preference for honesty. Finally, the desire to reveal the choice set can be based on a third motivation, that is neither related to others’ perception of the decision-maker’s niceness or the emotions of others: the desire for being honest to others. In the absence of revealing the choice set, observers will be left with an incorrect view of the decision maker’s choice, and if people simply care that the view is incorrect (without necessarily desiring that the view be positive), they might be driven to reveal the choice set. It is reasonable to consider this drive because honesty is a highly valued character trait (Anderson, 1968; Saxe, 1991), and strongly associated with the concept of morality and integrity in many cultures and religions (Carter, 1996). As such, it represents a construct that is fundamentally distinct from the two universal dimensions (warmth and competence) (Goodwin, 2015), and it has been shown that people even prioritize obtaining information about others’ honesty and morality, rather than learning about their competence or warmth (Brambilla, Rusconi, Sacchi, & Cherubini, 2011). The desire to appear honest can be interpreted as a special case of impression management, however, it is worth noting that being perceived as honest does not necessarily imply that the person will be liked, since the person can still be perceived as mean (“brutally honest”), thus we should treat the desire for honesty as separate from the desire to be perceived as nice. Most previous research on revealing behavior has, as mentioned before, used revealing primarily as a tool to detect concern for signaling generosity or hiding selfishness. In contrast, in the “lesser of two evils” situations, revealing could also be used simply to be or appear honest: While people who are primarily concerned about being perceived as fair or generous would not be willing to reveal their choice after selecting the *worse* option—they should only be willing to do so after selecting the *better* option—people who desire to be honest would reveal their choice set unconditionally, i.e., regardless of what their choice was. In the following sections, we turn to an empirical investigation of these motives and study revealing and hiding behavior in two behavioral experiments.

EXPERIMENT 1

THE LESSER OF TWO EVILS

To study the desire to reveal the choice set, we used an abstract laboratory study. We put decision-makers into a dictator game context that only allowed them to choose between two *bad* allocation choices, i.e., options that *both* favored the decision-maker over the recipient. Furthermore, they were told that recipients would not know the options the decision-makers faced, only the one option they chose. This asymmetry of information could lead to a fundamental attribution error: Allocators would initially appear selfish, regardless of their allocation choice, even if they chose the *more* generous option by sacrificing some of their money. After the chosen option was shown to the recipient, the allocators were offered the opportunity to reveal the choice set to the recipient for a price, and they were told that the recipient would know about this decision. Under this context, we investigated whether allocators were willing to sacrifice some of their money in order to have their intentions be understood—that is, whether they were willing to pay to disclose information that had not initially been available to the recipient.

Across three conditions we manipulated whether revealing the choice set signaled that the decision-maker's intention was positive (i.e., generous), neutral (reasonable), or negative (selfish). If people care primarily about either looking generous or about the feelings of others, they should be more likely to reveal the choice set in the positive and neutral conditions than in the negative condition. Furthermore, people could be motivated to *improve* their image or the feelings of others and separately to *avoid hurting* their image or the feelings of others. When people can reveal clear generosity (i.e., that they had sacrificed some of their own money to give their partner more), they get to achieve both goals. But when their choice was simply reasonable (i.e., it increased both their own and the other person's payoffs), this would not make them look especially good; it would only avoid hurting their image. This implies that people would be less willing to pay for disclosing information when doing so merely signals that they had been reasonable than if they could reveal their self-sacrificing generosity.

However, if people also care about honesty or being understood by others, we should see a significant amount of people revealing in the negative condition. Furthermore, if this latter drive represents the primary, or only, motivation to reveal, then revealing behavior should not differ across conditions.

Experimental procedure

We recruited 240 Amazon Mechanical Turk (AMT) workers (108 female, mean age = 35.5) to participate in a short study on decision-making. We report instructions and materials used in Experiment 1, as well as additional analyses and information about payments in Supplementary S-1. Participants were randomly assigned to the role of the allocator or the recipient.

In the first stage, every pair played a one-shot, anonymous, discrete choice dictator game with asymmetric information. The allocator had to choose between two possible allocations: “Option A” or “Option B” (see Table 1). The allocator was told that the recipient would know that she (i.e., the allocator) faced two options, but that the recipient would only see the option the allocator selected—the other option would remain unknown, though the recipient would be aware there was another option in the choice set. Thus, in the first stage, there was an information asymmetry between the allocator, who knew the full choice set, and the recipient, who knew only the chosen option. The recipient was told the same instructions, making her also aware of the information asymmetry. At this point, the allocator was *not* aware that she would later be given a chance to reveal the other option in the second stage.

Before making the allocation decision, participants were randomly assigned to one of three conditions: “Generous”, “Reasonable”, or “Selfish.” The conditions are named for how the allocator’s choice of Option A would most likely be perceived, given the other option. A novel feature of our design eliminates wealth effects across conditions: We designed the allocations such that the most frequently chosen option had the same payoffs across conditions. Thus, Option A is kept constant across conditions, while Option B varied. Option B was designed such that most participants would select Option A (\$1.80 / \$0.60), the social-welfare maximizing option, in all conditions. In the Generous condition, Option B meant giving \$2.00 to the allocator and \$0.20 to the recipient. We expected that most participants would select Option A because it offered them a cheap way to do something generous. In the Reasonable condition, Option A offered *both* people a gain of 40 cents relative to Option B (\$1.40 / \$0.20), and thus, it was the reasonable choice over B. In the Selfish condition, we made it relatively expensive to be generous, thus we expected that most participants would be attracted to the selfish Option A.

THE LESSER OF TWO EVILS

Table 1

Experimental conditions and payoffs, Experiment 1.

Condition ¹	Option A		Option B		Option A vs Option B ²
	Allocator	Recipient	Allocator	Recipient	
Generous (G)	\$1.80	\$0.60	\$2.00	\$0.20	-20¢ / +40¢
Reasonable (R)	\$1.80	\$0.60	\$1.40	\$0.20	+40¢ / +40¢
Selfish (S)	\$1.80	\$0.60	\$1.40	\$0.80	+40¢ / -20¢

¹ The names of conditions reflect the allocator's intention, had she chosen Option A.

² Change in payoffs (Allocator / Recipient) if Option A is selected over Option B.

Manipulating Option B across conditions allowed us to investigate whether allocators were more likely to reveal Option B when revealing signaled their generosity compared to when revealing highlighted something neutral, like reasonableness, or something negative, like selfishness.

One important assumption of this design is that the allocator's choice, Option A, looks selfish to the recipient who does not know the choice set. To ensure this was the case, we assessed whether the recipient was unhappy about the allocation, and whether she believed that the alternative, Option B, would have granted her a larger amount of money. After receiving the allocator's transfer, we asked the recipient to rate this transfer ("How do you feel about your partner's transfer?" -100: *very negative* ... 0: *neutral* ... +100: *very positive*) and elicited her expectations about the hidden alternative ("How much would Option B have given to your partner / to you?"). We used an incentive compatible method to elicit expectations, see [Appendix A1](#).

After the allocator reached a decision, we offered her an opportunity to reveal the unchosen alternative to the recipient. This opportunity was a surprise: The allocator was not told in advance that she would have the opportunity to reveal the other option. To reveal the unchosen alternative, she had to pay \$0.05 out of the money she previously allocated to herself. Keeping the alternative option hidden did not cost anything. (We included this cost to reduce spurious choices to reveal and to determine whether this behavior was valued enough to give up money to perform it.) After the allocator reached a decision, the recipient was informed about it, and was then asked to rate the allocator's transfer again. We hypothesized that these second ratings would be more favorable than the initial ratings, if and only if the revealed choice signaled the intention

THE LESSER OF TWO EVILS

to be generous or reasonable, but not selfish. In the absence of revealing, the recipient would not obtain any new information, so we did not expect ratings to change in those cases.

Results

Recipients' initial ratings and guesses. To confirm that recipients evaluated Option A negatively in isolation, first we looked at the initial guesses and ratings made by recipients *who received Option A* across all conditions ($n = 84$). The distribution of guesses about the alternative option (B) suggests that the experimental manipulation worked: Most recipients guessed that the allocator could have given them more money with Option B than Option A ($n = 67$, 80%). The majority also rated the initial transfer (Option A) negatively or neutrally ($n = 56$, 67%, see Table 3 for averages and CIs). Furthermore, an OLS linear regression analysis revealed a significant negative correlation between recipients' initial ratings of Option A and their guesses about Option B: Higher guesses about Option B are associated with lower ratings of Option A, $\beta = -24.4$, $t(83) = 2.441$, $p = .017$ (see Supplementary S-1).

Allocator's choice. The majority chose Option A in each condition (see Table 2). Consistent with both a concern about maintaining a positive image and for the feelings of others, among allocators who chose Option A, more people decided to pay for revealing their alternative option in the Generous (12 out of 23, 52%) and Reasonable (13 out of 39, 33%) conditions than in the Selfish condition (1 out of 22, 5%). Chi-square tests of goodness-of-fit yielded a significant difference and a medium to large effect size in the willingness to reveal between the Generous and Selfish conditions, $\chi^2(1, N = 45) = 12.417$, $p < .001$, $\phi = 0.525$, and between the Selfish and Reasonable conditions, $\chi^2(1, N = 61) = 6.592$, $p = .010$, $\phi = 0.329$, but not between the Generous and Reasonable conditions, $\chi^2(1, N = 62) = 2.134$, $p = .144$, $\phi = 0.186$. It is also worth mentioning that, among those who selected Option B in the Selfish condition—and thus, acted generously at a large cost—6 out of 18 allocators (33%) revealed Option A, providing further support for the hypothesis that allocators are more likely to reveal their choice set when it signals clear generosity. By contrast, only three people, across all conditions, paid to reveal the choice set when it made them look bad or unreasonable, suggesting little to no preference for being transparent or honest.

THE LESSER OF TWO EVILS

Table 2

Proportion of allocators choosing and revealing certain options across experimental conditions

Experiment	Condition ¹	<i>n</i>	Chose A		Revealed B ²		Chose B		Revealed A ³	
			<i>n</i>	%	<i>n</i>	%	<i>n</i>	%	<i>n</i>	%
Experiment 1	Reveal-G	40	23	58%	12	52%	17	43%	1	6%
	Reveal-R	40	39	98%	13	33%	1	3%	1	100%
	Reveal-S	40	22	55%	1	5%	18	45%	6	33%
Experiment 2	Reveal-G	81	60	74%	30	50%	21	26%	6	29%
	Reveal-S	81	60	74%	19	32%	21	26%	8	38%
	Hide-G	80	60	75%	55	92% ⁴	20	25%	14	70% ⁴
	Hide-S	86	62	72%	54	87% ⁴	24	28%	22	92% ⁴

¹R: "Reasonable"; G: "Generous"; S: "Selfish", ²Only among allocators who chose Option A, ³Only among allocators who chose Option B

⁴For the Hide conditions in Experiment 2, the table displays the proportion of allocators who kept the alternative option visible.

Recipients' final ratings. Finally, we look at recipients' ratings *after* the alternative option was revealed (or was kept hidden) to investigate whether these final ratings differ from the initial evaluations (see Table 3). Consistent with expectations, recipients' ratings of the transfer improved dramatically in those cases—and only in those cases—when revealing signaled the generosity or reasonableness of the allocator. We performed within-subject paired *t*-tests to determine whether these changes were significant and found that the increase in ratings was significant in all three conditions.¹ By contrast, ratings did not change significantly (all *ps* > .10) and stayed predominantly negative or neutral after allocators decided to keep hidden the alternative option.

¹ In the Selfish condition, we compared ratings of recipients whose partner initially selected Option B, and then, by revealing Option A, signaled their generosity.

THE LESSER OF TWO EVILS

Table 3

Recipients' ratings of the allocator's transfer before and after the allocator revealed (or did not reveal) the alternative option, Experiment 1.

Condition	<i>n</i>	Received option	<i>n</i>	Rating of transfer (before)		Allocator revealed	<i>n</i>	Rating of transfer (after)		Change of rating	
				<i>M</i>	95% CI			<i>M</i>	95% CI	<i>M</i>	sig. ²
Generous	40	A	23	-27.26	[-43.47, -11.06]	no	11	-17.82	[-42.72, 7.08]	5.64	.586
						yes	12	65.00	[38.73, 91.27]	95.75	< .001
		B	17	-82.76	[-95.5, -70.03]	no	16	-75.88	[-92.52, -59.23]	7.06	.146
						yes	1	[N/A] ¹			
Reasonable	40	A	39	-23.92	[-39.34, -8.5]	no	26	-28.38	[-46.97, -9.78]	-2.13	.677
						yes	13	45.69	[7.2, 84.18]	64.23	< .001
		B	1		[N/A] ¹	no	0	[N/A] ¹			
						yes	1	[N/A] ¹			
Selfish	40	A	22	-0.45	[-27.45, 26.54]	no	21	-14.57	[-38.8, 9.65]	-9.33	.430
						yes	1	[N/A] ¹			
		B	18	-6.39	[-32.33, 19.55]	no	12	7.42	[-20.71, 35.55]	1.67	.706
						yes	6	90.00	[75.46, 104.54]	120.67	.010

¹ Sample size was too small to calculate CI, ² Significance levels were obtained by performing paired sample *t*-tests on ratings of transfers.

EXPERIMENT 2

Experiment 1 provides an important proof of concept—that people are willing to pay to explain their choice, even *ex post*—and also suggests this behavior is *not* motivated by the desire to be or appear honest. Rather, the results suggest the behavior is driven either by allocators' desire to look generous (or to avoid looking selfish) or by their concern for others' feelings. However, Experiment 1 also has a potential confound: Although the design encourages most participants to select Option A regardless of condition, there is still a sizeable portion of people who select Option B. This means that there may be some self-selection such that the allocators who choose Option A in the Generous condition systematically differ from the allocators who choose Option A in the Selfish condition (e.g., the latter group may have stronger self-regarding preferences, and more reluctant to give up any money), which might serve as an alternative explanation for the observed difference in the allocator's behavior across conditions. The primary goal of Experiment 2 was to replicate the main results of Experiment 1—that allocators are willing to pay to reveal the choice set to explain their choice, especially when doing so reflects positively on them—but in a more controlled setting, ruling out potential self-selection biases.

THE LESSER OF TWO EVILS

Secondly, not all generously-behaving allocators decided to reveal their choice set, and we wanted to understand whether this was a fully-informed choice—implying that they did not care much about others’ beliefs—or whether it stemmed from misunderstanding the beliefs and feelings of the recipient. Did some allocators incorrectly believe that their partner was satisfied with the transfer, and for that reason find it unnecessary to reveal the choice set? Or did they know that their partner was dissatisfied, yet did not care enough about how their partner felt? To assess this, we included measures of allocators’ second-order beliefs (i.e., beliefs about the recipients’ beliefs).

Third, even if we do observe that people are more likely to reveal if they chose generously than if they chose selfishly, we will not be able to say that those revealers do not care about honesty at all—revealing signals both honesty and generosity in those cases. Thus, in Experiment 2 we also introduced new conditions in which the choice was to *hide* rather than to reveal the choice set. If decision-makers care primarily about looking generous or managing others’ feelings and do not care about being honest, then we should find that those whose choice appears selfish are willing to pay to hide the choice set. This would also mean that people would be willing to pay to *create* ambiguity in order to prevent a selfish choice from being known with certainty. The fact that many people in Experiment 1 left the choice set hidden suggests that many decision-makers did not have a problem with leaving another person in the dark about their actions. However, leaving the misunderstanding in place was a passive choice; it is not clear to what extent it was actively preferred. With the new conditions, we will be able to see whether people really do not care if another person misunderstands their decision, as long as they can protect either their image or the other person’s feelings.

Lastly, we included explicit measures to probe the underlying motivations for revealing (and hiding) as well as the effects of these actions on the recipients. The new measures assessed individual differences in allocators’ concern for others’ feelings, allocators’ desire to maintain a positive impression, allocators’ beliefs about the recipient’s curiosity, and recipients’ perceptions of allocator’s trustworthiness.

Experimental procedure

We recruited 674 AMT workers for Experiment 2. Because one of the main goals was to control for self-selection, we also wanted to rule out potential confounds due to high attrition

rates. As Zhou and Fishbach pointed out, high attrition rates are common in online studies, especially on AMT, and this can lead to major confounds and biases in the collected data (Zhou & Fishbach, 2016). However, we observed a very low attrition rate in Experiment 2, ruling out any potential confound due to attrition effects: Only 18 participants (2.7%) quit before completion. Six hundred fifty-six participants (293 female, mean age = 36.3) completed the experiment, and we included all of them in the final sample. We report the *a priori* power analysis and the pre-registration for Experiment 2 in Supplementary S-2.

Experiment 2 had a 2x2 between-subjects design: We manipulated whether the allocator's choice could be seen as generous or selfish after revealing the choice set, and whether the alternative option would be kept hidden by default (and could be revealed at a cost, as in Experiment 1), or would be revealed by default (and could be hidden at a cost).

The experimental procedure in the Reveal conditions was similar to the conditions in Experiment 1, with the following differences: We removed the Reasonable payoff set, and thus, used only the Generous and Selfish payoff sets. Then, we sorted participants randomly into these two conditions *after* they chose either Option A or B. To avoid any deception, we did this in the following way. First, the allocators were told that they had to choose between two options, but they had only probabilistic information about one option. Option A was always the same as in Experiment 1 (\$1.80 for the allocator and \$0.60 for the recipient), while Option B could be either \$2.00 / \$0.20 or \$1.40 / \$0.80 (the same as in the Generous and Selfish conditions in Experiment 1). Allocators were told that Option B was one of these allocations, but at the time of their choice, they were not able to tell which one (they were told that each allocation for Option B had a 50% chance to be realized, and the actual allocation would be revealed immediately after their choice). After allocators made their choice, we randomly sorted them into two conditions: In the Generous condition Option B was \$2.00 / \$0.20, while in the Selfish condition Option B was \$1.40 / \$0.80.

The second major difference from Experiment 1 was that we also elicited allocators' second-order beliefs about recipients' guesses and ratings: Each allocator was incentivized to predict their partner's guess about the alternative option in the first stage and how their partner would rate their transfer in the second stage (after the alternative option had been or had not been revealed). By comparing recipients' first-order beliefs and guesses to these new second-order measures, we can assess whether non-revealing behavior is driven by inaccurate second-order

THE LESSER OF TWO EVILS

beliefs (e.g., non-revealers believe that the recipients rate the transfer positively). Furthermore, as one might argue, the wording of the rating used in Experiment 1 (“How do you feel about your partner’s transfer?”) might have been slightly misinterpreted by some recipients, who, instead of rating the allocator’s *act* of choosing a certain allocation and transferring some amount, might have simply rated the transfer (the amount of money received) itself. To rule out any such misinterpretation, we asked recipients a second question both before and after revealing: “How do you feel about your partner as a person?” *-100: Very untrustworthy ... 0: Neither untrustworthy, nor trustworthy ... +100: Very trustworthy*. This question ought to measure recipients’ attitude towards the allocators. We also asked allocators to indicate their second-order beliefs about their partner’s trustworthiness ratings. Note that if participants interpret the transfer rating question correctly (measuring attitude towards the allocator and her actions), then both the transfer ratings and trustworthiness ratings measure the same construct, thus we should expect a very strong and positive correlation between. However, if there is any misinterpretation of the transfer ratings, then we should expect a divergence between transfer and trustworthiness ratings, as they do not necessarily measure the same constructs.

Third, we asked recipients in the first stage how curious they were about the alternative allocation (“To what extent do you want to know the unchosen option?” *0: I do not want to know at all ... 100: I strongly want to know*), and we elicited allocators’ beliefs about their partner’s curiosity about the unselected option. Strong feelings of curiosity are often associated with negative emotions, feelings of deprivation, and people find such situations aversive (Litman & Jimerson, 2004; Loewenstein, 1994). Thus, if allocators believed that their partner was very curious about the (hidden) unchosen option, they could have decided to reveal the choice set to alleviate such negative feelings. Therefore, this measure about their partner’s curiosity ought to capture such concerns for recipients’ emotions.

To assess differences in social preferences as a potential explanation for the decision to reveal, we included six questions after the second stage to investigate whether allocators care about the feelings of strangers and close others, and whether they want to be perceived as fair and honest by strangers and close others (see [Appendix A2](#)). Since the experiment featured an anonymous one-shot interaction, we included questions about close others to serve as control measures, which should not have any association with the allocators’ decision to reveal. By contrast, we expected a significant correlation between allocators’ concerns for strangers (and the

THE LESSER OF TWO EVILS

desire to be perceived fair and honest by them) and the decision to reveal, moderated by the experimental condition: Allocators who have stronger concerns for strangers would be more likely to reveal the choice set in the Generous condition, but less likely to reveal the choice in the Selfish condition.

The procedure in the Hide conditions was the same as in the Reveal conditions described above, but the allocators were told that the unchosen alternative would be revealed to the recipients by default. Then, they had the costly opportunity to hide the alternative option for 5 cents. If they decided to hide the alternative, the recipient was informed about this, and was shown only the chosen option. Everything else, including the first- and second-order belief measures was the same as in the Reveal conditions. Thus, the Hide conditions mirrored the Reveal conditions: Whereas in the Reveal conditions the allocators could pay to reveal the unchosen alternative, in the Hide conditions they could pay to hide it.

Results

Allocator's choice. The majority of allocators selected Option A (74%). We limit all subsequent analyses to these participants. Consistent with our previous findings and basic impression management goals (i.e., want to look positive), allocators were more willing to reveal their alternative if they could signal their generosity and/or make the other person feel good, even after controlling for self-selection. A Chi-square tests of goodness-of-fit yielded a significant difference in the willingness to reveal between the Reveal-Generous (30 out of 60, 50%, see Table 2) and Reveal-Selfish (19 out of 60, 32%) conditions, $\chi^2(1, N = 120) = 4.174, p = .041, \phi = 0.186$. This finding rejects the hypothesis that the difference in revealing behavior between Generous and Selfish conditions can be explained *entirely* by self-selection. However, the much smaller effect size obtained in Experiment 2 ($\phi = 0.186$, cf. $\phi = 0.525$ in Experiment 1) suggests that in the first experiment the large differences were *partially* caused by self-selection. More interesting is, however, the finding that a substantial proportion of allocators (32%) were willing to pay to reveal the choice set that made the recipient think that they had been selfish. This is inconsistent with wanting to look generous or wanting to make the other person feel good, and it is more consistent with the idea that a substantial portion of people do care about honesty in and of itself, even when it signals selfishness and/or hurts the feelings of others. Participant comments further support this idea, as mentioned later in the General Discussion.

THE LESSER OF TWO EVILS

By contrast, almost no one was willing to pay to hide their choice set (13 out of 122, 11%), and a Fisher's exact test revealed no significant difference in the willingness to hide between the Hide-Generous (5 out of 60, 8%) and Hide-Selfish (8 out of 62, 13%) conditions, $p = .560$. A direct comparison between the Reveal-Selfish and Hide-Selfish conditions further substantiates this claim: A Fisher's exact test revealed that allocators were significantly more likely to pay to reveal their selfishness (32%), than to hide it (13%), $p = .016$. This asymmetry between the Reveal and Hide conditions (i.e., lot of people pay to reveal their choice but very few people pay to hide their choice set) suggests that people in this scenario do not care enough about looking generous or enough about the other person's feelings that they are willing to obfuscate their choice set and create a state of misunderstanding in the other person in order to obtain that goal.

*Accuracy of allocator's second-order beliefs.*² One simple explanation for why 48% of allocators failed to reveal the choice set even when it could have put them in a good light, is that these people made a genuine mistake, informed by their biased beliefs: They (incorrectly) believed that their partner would be satisfied with the transfer, even without knowing what the alternative was. By comparing allocators' second-order beliefs about the recipients' ratings to recipients' actual ratings we can investigate whether there were any systematic biases in allocators' beliefs.

² The results discussed in this section are based on transfer ratings only. We found a significant and strong positive correlation between transfer and trustworthiness ratings, suggesting that the two ratings are measuring a similar underlying construct. A factor analysis confirmed this, and results clearly indicated the existence of a single component. Therefore, we report results for one measure only to avoid redundancy. Details of this factor analysis are reported in Supplementary S-2.

THE LESSER OF TWO EVILS

Figure 1. Allocators' and recipients' mean rating of the transfer (Option A) after the alternative (Option B) was revealed (or kept hidden). Error bars indicate ± 1 SE.

Figure 1 illustrates the first- and second-order beliefs of participants after the alternative option had been revealed (or kept hidden). First, when looking at recipients' ratings (light grey bars), we see that the ratings on average were predominantly negative or neutral, except for when *generous* allocators revealed their choice set or kept it revealed, in which cases the average rating was positive, replicating the results of Experiment 1. Independent samples *t*-tests revealed that recipients' ratings in the Reveal-Generous condition were significantly higher after the alternative had been revealed ($M = 63.7$) than after the alternative had been kept hidden ($M = -26.4$), $t(58) = 6.549$, $p < .001$, Cohen's $d = 1.691$. Although recipients in the Reveal-Selfish condition evaluated the transfer negatively regardless whether the alternative had been revealed ($M = -15.40$) or kept hidden ($M = -40.17$), revealing was associated with marginally significantly better ratings, $t(58) = 1.846$, $p = .070$, Cohen's $d = 0.504$. A similar pattern emerged in the Hide conditions: Recipients in the Hide-Generous rated the transfer significantly higher after the alternative was left visible ($M = 45.93$) than after the alternative had been hidden ($M = -15.40$), $t(58) = 2.378$, $p = .021$, Cohen's $d = 1.226$. Recipients in the Hide-Selfish condition also rated the transfer better after the choice set was left visible ($M = 1.98$) than after it had been hidden (M

THE LESSER OF TWO EVILS

= -25.38), however, this difference was not significant, $t(58) = 1.377$, $p = .174$, Cohen's $d = 0.492$.

Next, we looked at allocators' second-order beliefs about recipients' ratings. If allocators who did not reveal the alternative in the Reveal-Generous condition had biased beliefs, then we should expect them to hold overly optimistic beliefs about the recipients' ratings. However, we find that allocators were surprisingly accurate in the Reveal-Generous condition. Those who revealed the alternative believed that their transfer would be rated very positively ($M = 61.3$), whereas those who did not reveal correctly inferred that their transfer would be seen negatively ($M = -14.4$). The difference between the beliefs of revealers and non-revealers is large and significant, $t(58) = 6.316$, $p < .001$, Cohen's $d = 1.632$. Furthermore, these second-order guesses about recipients' ratings were not significantly different from how recipients actually rated the transfer, $t(58) = 0.194$, $p = .847$, Cohen's $d = 0.050$, and $t(58) = 0.895$, $p = .375$, Cohen's $d = 0.231$, in cases when the choice set had been revealed and had been kept hidden, respectively. Thus, allocators in the Reveal-Generous condition anticipated correctly that their transfer would be rated negatively if they did not reveal the choice set. This finding clearly rules out that the decision to not reveal was based on biased beliefs.

By contrast, allocators in the Reveal-Selfish condition tended to systematically and significantly overestimate recipients' ratings, regardless of whether they had revealed the choice set, $t(118) = 2.831$, $p = .005$, Cohen's $d = 0.517$. However, allocators' guesses were still directionally consistent with recipients' actual ratings: Allocators who revealed thought that their transfer would be rated significantly higher ($M = 16.3$) than allocators who did not reveal ($M = -16.9$), $t(58) = 2.446$, $p = .017$, Cohen's $d = 0.662$. This pattern is consistent with how recipients evaluated the transfer in the Reveal-Selfish condition, and the similar effect sizes ($d = 0.662$ and $d = 0.504$, for allocators and recipients, respectively) suggest that allocators had accurate beliefs about the magnitude of the difference in the ratings between cases when the choice set had been revealed and cases when it had been kept hidden. Again, these results indicate that allocator's behavior whether to reveal the alternative could not be primarily driven by biases in their beliefs about what their partners would think or feel—most non-revealers were aware that their transfer would be viewed negatively if they failed to reveal their choice set.

Allocators also held accurate beliefs in the Hide-Generous condition. Allocators' belief about the recipients' rating after the alternative had been kept revealed ($M = 38.60$) was not

THE LESSER OF TWO EVILS

significantly different from recipients' actual ratings ($M = 45.93$), $t(108) = 0.708$, $p = .481$, Cohen's $d = 0.135$. Allocators were slightly pessimistic in the Hide-Selfish condition: Allocators' belief about the recipients' rating after the alternative had been kept revealed ($M = -17.13$) was marginally significantly lower than how recipients actually rated the transfer in these cases ($M = 1.98$), $t(106) = 1.955$, $p = .053$, Cohen's $d = 0.376$. We did not conduct any t -tests measuring the accuracy of allocators' beliefs in cases when the alternative had been hidden in the Hide conditions, because the small sample sizes in these cases ($n = 5$ and $n = 8$, in the Hide-Generous and Hide-Selfish, respectively) would not warrant that the basic assumptions for t -tests would hold.

Individual predictors of the decision to reveal. Since the vast majority of allocators (89%) in the Hide conditions kept the alternative option visible, the sample of allocators who hid the alternative option in these conditions was too small to be able to draw any reliable statistical inference about their individual characteristics, preferences, or beliefs. Thus, we limit the analyses in the following section to allocators who participated in the two Reveal conditions ($n = 120$). Note also, that because the number of people who hid the alternative in the Hide conditions was small (13/122), we do not run similar regressions on the Hide conditions.

First, we examine the measures assessing participants' concerns towards the beliefs and feelings of strangers and close others. We observed significant positive correlations between all six measures, all r s $> .279$, all p s $< .001$, suggesting reasonable factorability. Thus, we conducted a principal components analysis and identified two primary factors that could account for these six measures: "attitude towards close others" and "attitude towards strangers" (for detailed results, see [Appendix A2](#)). These compound measurements indicate to what extent allocators cared about the beliefs and feelings of close others or strangers, respectively. The compound measures were created by taking the mean of the three corresponding items, and had reasonable reliability, Cronbach's $\alpha = 0.77$ (close others) and $\alpha = 0.89$ (strangers). We then conducted an OLS linear regression analysis, with allocators' decision to reveal the alternative (or keep it hidden) as the dependent variable, and the following variables as potential predictors: condition (Generous or Selfish), attitude towards close others, attitude towards strangers, second-order belief about recipient's guess about recipient's alternative payoff, second-order belief about recipient's curiosity, and demographic factors (sex, age). To study the effects of different predictors, we created six models, gradually adding new predictors, or omitting non-predictive

THE LESSER OF TWO EVILS

variables. We also considered possible interactions between condition and the other predictors. Table 4 summarizes the results of the regression analysis and the resulting six models.

Table 4
OLS linear regressions: allocators' decision to reveal, Experiment 2

	Model I	Model II	Model III	Model IV	Model V	Model VI
Constant	0.500* (0.063)	0.011 (0.363)	0.131 (0.188)	-0.132 (0.225)	-0.145 (0.247)	-0.024 (0.285)
Condition: Selfish	-0.183* (0.089)	0.536 (0.551)	0.291 (0.265)	0.607 [†] (0.324)	0.386 (0.402)	0.471 (0.432)
Attitude towards close others	—	0.006 (0.004)	—	—	—	—
Attitude towards close others X Condition: Selfish	—	-0.008 (0.006)	—	—	—	—
Attitude towards strangers	—	—	0.006* (0.003)	0.006* (0.003)	0.006* (0.003)	0.006* (0.003)
Attitude towards strangers X Condition: Selfish	—	—	-0.007 [†] (0.004)	-0.008* (0.004)	-0.008* (0.004)	-0.008* (0.004)
2nd order belief about recipient's alternative payoff	—	—	—	0.258* (0.125)	0.253 [†] (0.132)	0.223 [†] (0.129)
2nd order belief about recipient's alternative payoff X Condition: Selfish	—	—	—	-0.305 [†] (0.176)	-0.297 (0.181)	-0.270 (0.180)
2nd order belief about recipient's curiosity	—	—	—	—	0.000 (0.003)	—
2nd order belief about recipient's curiosity X Condition: Selfish	—	—	—	—	0.003 (0.004)	—
Sex: female	—	—	—	—	—	0.187 (0.140)
Sex: female X Condition: Selfish	—	—	—	—	—	-0.218 (0.190)
Age	—	—	—	—	—	-0.003 (0.006)
Age X Condition: Selfish	—	—	—	—	—	0.004 (0.009)
Adjusted R^2	.027	.028	.048	.068	.061	.050
Model Significance	.041	.098	.033	.023	.048	.097

Note: the OLS linear regression analyses above are limited to allocators who chose Option A in Reveal conditions ($n = 120$). Coefficients are unstandardized coefficients with standard errors in parentheses. Dependent variable is the choice whether to reveal Option B (1) or keep it hidden (0).

* $p < .05$, [†] $p < .10$, two-tailed.

Model I has a single predictor (condition) and yields the same result as the Chi-square test reported earlier in this section: Allocators are significantly less likely to reveal the alternative if they are in the Selfish condition. In Models II and III we add the two compound variables (attitude towards close others and attitude towards strangers) and the interaction terms between these variables and condition. As expected, the attitude towards close others does not predict the

THE LESSER OF TWO EVILS

decision to reveal, and adding this variable as a predictor to Model II does not improve the predictive power. Thus, we omit this variable from subsequent models.

By contrast, the attitude towards strangers is a significant predictor of the decision to reveal, and improves the predictive power of the model. Model III indicates that people who care more about strangers' feelings and how strangers perceived them are more likely to pay to reveal the alternative option. Also, a marginally significant interaction emerges between this variable and the condition, which suggests that in the Selfish condition, allocators who care more about strangers are *less* likely to reveal the alternative. Note that adding the main effect of attitude and the interaction between attitude and condition makes the main effect of condition disappear, which suggests that the main difference between the Generous and Selfish conditions is primarily driven by the differential behavior of people who care about the attitude of strangers.

So far, we have established that individuals' concern for strangers predicts whether they will reveal the choice set to anonymous strangers or keep it hidden. What about allocators' second-order beliefs about recipients' expectations? Surely, it feels much worse to withhold the information that the alternative option would have been terrible for the recipient, when she was expecting the opposite, than when she was expecting that the alternative was indeed terrible. But are allocators more likely to reveal the alternative, when they believe that their partner thinks that the alternative option would have been more favorable to them? To answer this question, we add allocators' second-order guesses about the recipients' alternative payoff to the linear regression (Model IV). The results clearly support that allocators' beliefs about recipients' expectations matter: Allocators who thought that their partner expected a better alternative were more likely to reveal the choice set in the Generous condition (when the alternative was in fact worse for the recipient), but *less* likely to reveal it in the Selfish condition (when the alternative was in fact better). Adding these second-order beliefs about recipients' expectations increases the predictive power of the model, and more importantly, the attitudes towards strangers remain a significant predictor. This means that both attitudes toward strangers *and* beliefs about expectations predicts allocators' decision to reveal the alternative.

What about curiosity? Were allocators more likely to reveal if they believed that their partner was more curious, thereby satisfying their curiosity? In Model V we added allocators' second-order beliefs about their partner's curiosity. However, neither this variable, nor its interaction

THE LESSER OF TWO EVILS

with the condition predicts the allocators' action. Thus, allocators do not reveal their choice set to merely satiate other's curiosity.

Finally, in Model IV we test whether adding demographic factors (sex and age), and their interactions with condition can improve the predictive power of the model. However, neither variables are significant predictors, thus there is no evidence that demographic factors per se have any predictive power in explaining the decision to reveal the alternative. In sum, Model IV emerges as the model with the highest predictive power and suggests that attitudes towards strangers and beliefs about recipients' expectations are the two factors that significantly predict the decision to reveal actions. The more an allocator cares about strangers, and the higher her guess about her partner's belief about the alternative is, the more likely she is to reveal her choice set in Generous condition, and the less likely she is to reveal her choice set in the Selfish condition.

GENERAL DISCUSSION AND CONCLUSION

When facing only bad options, one's choice is bound to create a misunderstanding, if the recipient of the choice or some other audience does not know the full set of choices, or is unaware of choice-limiting constraints. In this paper we investigated for what intrinsic reasons and to what extent people might be willing to reveal their choice set in such cases. We found that, many people are indeed willing to pay for revealing the context of their choice *ex post*, even to anonymous individuals in one-shot interactions. In Experiment 1, we demonstrated that, when the choice set reflected positively on the allocator, over half of allocators were willing to pay to reveal the choice set (52%). Fewer participants were willing to pay to reveal the other choice (33%) when revealing neither signaled generosity nor selfishness—only reasonableness—while almost no participants chose to reveal the choice set when it signaled selfishness (5%). In Experiment 2, we replicated the results of Experiment 1, and demonstrated that the main effects observed in Experiment 1 cannot be fully explained by self-selection of participants: More people preferred to reveal when it signaled generosity (“generous revealing,” 50%) than when it signaled selfishness (“selfish revealing,” 32%). However, this also made it clear that a significant portion of people reveal when it is likely to have a negative impact on both the allocator's image and on the recipient's feelings. This behavior of revealing selfish intentions seems largely to be

THE LESSER OF TWO EVILS

explained by a preference for either being or appearing transparent. In their freeform text comments, many of these allocators expressed the desire to be honest: e.g., “I felt like I should be honest about the reality of the decision, even though it showed I kept more for myself than for them”; “so my partner can see how everything really happened”; “I have nothing to hide!!” Others anticipated that the recipient would see their choice as reasonable or understandable: e.g., “They will see that I was just trying to maximize my bonus”; “I just thought it would make more sense to my partner if they saw what the options were, and it would explain why they received a lesser amount than my portion.”

Consistent with this, we found that revealing generous intentions, unlike selfish intentions, to anonymous recipients was correlated with a stronger concern for the thoughts and feelings of strangers (but not close others). The difference between Experiments 1 and 2 due to self-selection is potentially interesting in that it suggests that the people who tend to behave generously also tend to desire to be or appear honest. However, when such people were put in a situation in which they appeared selfish, and their actions could not be interpreted as generous no matter what they did, they still preferred to be honest by revealing their choice set.

We also found a remarkable asymmetry between the desire to reveal, and the desire to hide the choice set: While most people were willing to effectively hide their selfish intentions by *not revealing* the choice set, almost no one was willing to hide their selfish intentions when they had the opportunity to actively *hide* the choice set. That is, allocators seemed to be quite likely to passively allow a misunderstanding when it could help their image or prevent negative feelings in recipients, but they seemed decisively unwilling to actively create a misunderstanding in order to accomplish the same goal.

Did revealing behavior achieve its goal? We verified that this act of disclosing information was indeed effective at improving recipients’ ratings of the transfer and the perceived trustworthiness of the allocator in the generous conditions. In Experiment 2, we ruled out that non-revealing allocators underestimated these beneficial effects of revealing, and thus made a mistake by choosing not to reveal. On the contrary, non-revealers were aware that recipients expected a much better alternative option (before it was revealed) and that they were dissatisfied. These results provide strong support against the explanation that non-revealing behavior is driven by inaccurate beliefs, like those stemming from the curse of knowledge (Camerer, Loewenstein,

& Weber, 1989). Instead, as we found in Experiment 2, non-revealers had weaker concern about the thoughts and beliefs of strangers than revealers did.

If people were willing to pay to reveal the context when it would project good intentions, why were they not willing to pay to hide the context when it could hide negative intentions? Most people in the Hide-Selfish condition left their choices known to the recipient. One explanation is that, in this experimental setup, hiding the unchosen allocation would still project bad—although more ambiguous—intentions since both allocations asymmetrically favored the allocator. As one allocator commented, “Either way, I look greedy. I’d rather they just know the details.”

However, the act of hiding may also be qualitatively different than the act of revealing, signaling a negative trait: dishonesty. Indeed, of the allocators in the Hide-Selfish condition who did not pay to hide their choice set, a large portion (24 out of 54) explained this choice as the desire to be honest and transparent with their partner. By contrast, those who chose not to pay to reveal in the Reveal-Selfish condition in Experiment 2 explained their decisions in the comments primarily by saying revealing was too expensive (19 out of 41, 46%) or that they did not want their partner to know because it would make them look bad (16 out of 41, 39%). The difference is that keeping the choice set hidden required *action* in the Hide-Selfish condition but *inaction* in the Reveal-Selfish condition. As some research suggests, acts of commission are seen as more blameworthy than acts of omission (Spranca, Minsk, & Baron, 1991), and this may serve as an additional reason for why allocators were unwilling to pay to hide their choices.

Overall, these experiments have demonstrated that people are indeed willing to incur a cost to correct a misunderstanding by revealing their choice set even in a one-shot, anonymous setting, when there is no chance of future interaction, and there can be no reputational consequences. And they do this for a variety of reasons, with the strongest motivation deriving from wanting to send a positive signal, whether that is for self-serving image concerns or out of concern for the feelings of others. However, there are certainly some limitations and open questions. For instance, we do not have conclusive evidence regarding whether transparency and the act of revealing per se—regardless of the content to be revealed—has intrinsic value and is appreciated by the audience. We do not find an unconditional significant increase in ratings after the choice set is revealed—only if doing so reveals the generosity of the decision-maker. Therefore, it is possible that recipients did not appreciate the decision-makers’ honesty at all, rather, viewed these actions negatively (“brutal honesty”). By contrast, allocators seemed to think revealing

their selfishness would be perceived positively on average, suggesting those allocators may have misguided beliefs. Thus, it is not clear whether these allocators revealed because they thought honesty would make them look good to others or whether they are simply the type of people who want to live by the principle of honesty, regardless of how it affects others.

Furthermore, we did not manipulate the details of *how* and *when* the context would be revealed. In our experiments, it was always the decision-maker who had to make an active choice between revealing or withholding information, immediately following his or her decision. But in real life, often it is the receiver or a third-party (e.g., a journalist, a co-worker) who can reveal such information, or the information is discovered accidentally. Voluntary disclosure by the decision-maker can be in some cases associated with honesty and transparency, while in other cases it can raise suspicion about the decision-maker's strategic motives, or even be interpreted as bragging (Berman, Levine, Barasch, & Small, 2015). There can also be a considerable delay between the choice and revealing the context, which might attenuate or enhance the effect of disclosure. Finally, in our experiments the decision-makers could reveal the full context of the choice (i.e., all alternatives), and if they did so, the audience was aware that the full context had been revealed. By contrast, if the choice set is more complex and contains several alternatives, the decision-maker might be motivated to reveal only partial information. Similarly, when the audience can not verify whether the full context or only some part of it has been revealed, they might perceive information disclosure as less credible, and remain suspicious about the decision-maker's intentions. We believe that the above are all highly relevant questions, which have important implications for policy, public relations, and organizational behavior, and that the current paper serves as a solid basis for future research addressing these issues.

REFERENCES

- Anderson, N. H. (1968). Likableness ratings of 555 personality-trait words. *Journal of Personality and Social Psychology*, 9(3), 272–279. <https://doi.org/10.1037/h0025907>
- Andreoni, J., & Rao, J. M. (2011). The power of asking: How communication affects selfishness, empathy, and altruism. *Journal of Public Economics*, 95(7–8), 513–520. <https://doi.org/10.1016/j.jpubeco.2010.12.008>
- Andreoni, J., Rao, J. M., & Trachtman, H. (2017). Avoiding the Ask: A Field Experiment on Altruism, Empathy, and Charitable Giving. *Journal of Political Economy*, 125(3), 625–653. <https://doi.org/10.1086/691703>
- Battigalli, P., & Dufwenberg, M. (2007). Guilt in Games. *American Economic Review*, 97(2), 170–176. <https://doi.org/10.1257/aer.97.2.170>
- Berman, J. Z., Levine, E. E., Barasch, A., & Small, D. A. (2015). The Braggart’s Dilemma: On the Social Rewards and Penalties of Advertising Prosocial Behavior. *Journal of Marketing Research*, 52(1), 90–104. <https://doi.org/10.1509/jmr.14.0002>
- Bierbrauer, G. (1979). Why did he do it? attribution of obedience and the phenomenon of dispositional bias. *European Journal of Social Psychology*, 9(1), 67–84. <https://doi.org/10.1002/ejsp.2420090106>
- Brambilla, M., Rusconi, P., Sacchi, S., & Cherubini, P. (2011). Looking for honesty: The primary role of morality (vs. sociability and competence) in information gathering. *European Journal of Social Psychology*, 41(2), 135–143. <https://doi.org/10.1002/ejsp.744>
- Camerer, C., Loewenstein, G., & Weber, M. (1989). The Curse of Knowledge in Economic Settings: An Experimental Analysis. *Journal of Political Economy*, 97(5), 1232–1254. <https://doi.org/10.1086/261651>
- Cappelen, A. W., Halvorsen, T., Sørensen, E. Ø., & Tungodden, B. (2017). Face-saving or fair-minded: What motivates moral behavior? *Journal of the European Economic Association*, 15(3), 540–557. <https://doi.org/10.1093/jeaa/jvw014>
- Carter, S. L. (1996). *Integrity*. New York: Basic Books. Retrieved from <https://repository.library.georgetown.edu/handle/10822/1037599>
- Charness, G., & Dufwenberg, M. (2006). Promises and Partnership. *Econometrica*, 74(6), 1579–1601. <https://doi.org/10.1111/j.1468-0262.2006.00719.x>
- Dana, J., Cain, D. M., & Dawes, R. M. (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. <https://doi.org/10.1016/j.obhdp.2005.10.001>
- Dana, J., Weber, R. a., & Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1), 67–80. <https://doi.org/10.1007/s00199-006-0153-z>
- Dufwenberg, M., & Gneezy, U. (2000). Measuring Beliefs in an Experimental Lost Wallet Game. *Games and Economic Behavior*, 30(2), 163–182. <https://doi.org/10.1006/game.1999.0715>
- Fabrizio, T. (2016, March 27). “Lesser of Two Evils” Takes On New Meaning in 2016. Retrieved from https://www.realclearpolitics.com/articles/2016/03/27/lesser_of_two_evils_takes_on_new_meaning_in_2016_130104.html
- Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: warmth and competence. *Trends in Cognitive Sciences*, 11(2), 77–83. <https://doi.org/10.1016/j.tics.2006.11.005>
- Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin*, 117(1), 21–38. <https://doi.org/10.1037/0033-2909.117.1.21>
- Golman, R. (2016). Good manners: signaling social preferences. *Theory and Decision*, 81(1), 73–88. <https://doi.org/10.1007/s11238-015-9527-7>
- Goodwin, G. P. (2015). Moral Character in Person Perception. *Current Directions in Psychological Science*, 24(1), 38–44. <https://doi.org/10.1177/0963721414550709>
- Grossman, Z., & van der Weele, J. J. (2017). Self-Image and Willful Ignorance in Social Decisions. *Journal of the European Economic Association*, 15(1), 173–217. <https://doi.org/10.1093/jeea/jvw001>
- Heider, F., & Simmel, M. (1944). An Experimental Study of Apparent Behavior. *The American Journal of Psychology*, 57(2), 243. <https://doi.org/10.2307/1416950>
- Ho, B. (2012). Apologies as Signals: With Evidence from a Trust Game. *Management Science*, 58(1), 141–158. <https://doi.org/10.1287/mnsc.1110.1410>
- Jones, E. E., & Davis, K. E. (1965). From Acts To Dispositions The Attribution Process In Person Perception. In *Advances in Experimental Social Psychology* (pp. 219–266). [https://doi.org/10.1016/S0065-2601\(08\)60107-0](https://doi.org/10.1016/S0065-2601(08)60107-0)
- Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social*

- Psychology*, 3(1), 1–24. [https://doi.org/10.1016/0022-1031\(67\)90034-0](https://doi.org/10.1016/0022-1031(67)90034-0)
- Litman, J. A., & Jimerson, T. L. (2004). The Measurement of Curiosity As a Feeling of Deprivation. *Journal of Personality Assessment*, 82(2), 147–157. https://doi.org/10.1207/s15327752jpa8202_3
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116(1), 75–98. <https://doi.org/10.1037/0033-2909.116.1.75>
- Miller, D. T., & Nelson, L. D. (2002). Seeing approach motivation in the avoidance behavior of others: Implications for an understanding of pluralistic ignorance. *Journal of Personality and Social Psychology*, 83(5), 1066–1075. <https://doi.org/10.1037//0022-3514.83.5.1066>
- Mischel, W. (1977). On the future of personality measurement. *American Psychologist*, 32(4), 246–254. <https://doi.org/10.1037/0003-066X.32.4.246>
- Nisbett, R. E., & Ross, L. (1991). *The person and the situation*. New York: McGraw-Hill.
- Quattrone, G. A. (1982). Overattribution and unit formation: When behavior engulfs the person. *Journal of Personality and Social Psychology*, 42(4), 593–607. <https://doi.org/10.1037//0022-3514.42.4.593>
- Saxe, L. (1991). Lying: Thoughts of an applied social psychologist. *American Psychologist*, 46(4), 409–415. <https://doi.org/10.1037/0003-066X.46.4.409>
- Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, 27(1), 76–105. [https://doi.org/10.1016/0022-1031\(91\)90011-T](https://doi.org/10.1016/0022-1031(91)90011-T)
- te Velde, V. L. (2018). Beliefs-based altruism as an alternative explanation for social signaling behaviors. *Journal of Economic Behavior & Organization*, 152, 177–191. <https://doi.org/10.1016/j.jebo.2018.06.011>
- Trope, Y., & Liberman, A. (1993). The Use of Trait Conceptions to Identify Other People's Behavior and to Draw Inferences about their Personalities. *Personality and Social Psychology Bulletin*, 19(5), 553–562. <https://doi.org/10.1177/0146167293195007>
- Zhou, H., & Fishbach, A. (2016). The pitfall of experimenting on the web: How unattended selective attrition leads to surprising (yet false) research conclusions. *Journal of Personality and Social Psychology*, 111(4), 493–504. <https://doi.org/10.1037/pspa0000056>

APPENDIX

Appendix A1. Belief elicitation methods

Recipients were told that they had to make a guess about both parties' alternative payoff (i.e. their own and their partner's). They were also told that they had a chance to win a bonus of \$0.25 per guess (\$0.50 in total) if they guessed the unchosen alternative payoffs correctly. The probability of getting these bonuses increased with accuracy: Within an error of $\pm\$0.05$, they received the bonus with 100% probability ($\pm\$0.10$: 75%, $\pm\$0.20$: 50%, $\pm\$0.30$: 25%). If the difference between the guess and the actual value was higher than \$0.30, participants did not have a chance to win the bonus. After participants made their guesses, we calculated the probability of winning the bonus based on their accuracy, and then a random number generator algorithm determined whether the recipient has won \$0.50 (for correctly guessing both payoffs) or \$0.25 (for correctly guessing one payoff) or \$0. Recipients were not informed directly about their accuracy, or probability to win, but they were told whether they received any bonuses at the end of the experiment. This belief elicitation technique may seem overly complicated, but it was necessary in our design. We had to make sure that recipients could not determine the unchosen alternative option based on the feedback from the guessing stage. Disclosing their accuracy or probability to win the prize would have indirectly revealed the alternative option, even if the allocator decided to keep it hidden.

**Appendix A2. Additional dependent measures and Principal Component Analysis,
Experiment 2**

The following were asked of allocators *after* the choice to reveal or not had already been made.

1. “In your option, to what extent does your partner want to know Option [unselected]?
[slider: 0 does not want to know at all ... 100 strongly wants to know]

“Indicate response for the questions below [slider: 0 not important at all ... 100 very important]:

2. “How important is it for you that people who are close to you (family, friends, colleagues) perceive you as fair?” [fairness / close]
3. “How important is it for you that people who are close to you (family, friends, colleagues) perceive you as honest?” [honesty / close]
4. “How important is it for you to make sure that people who are close to you (family, friends, colleagues) don't feel bad?” [empathy / close]
5. “How important is it for you that people who you do not know (strangers) perceive you as fair?” [fairness / stranger]
6. “How important is it for you that people who you do not know (strangers) perceive you as honest?” [honesty / stranger]
7. “How important is it for you that people who you do not know (strangers) don't feel bad?” [empathy / stranger]”

We observed significant positive correlations between all of the six measures that captured allocator's social concerns towards strangers and close others, all $r_s > .279$, all $p_s < .001$, suggesting reasonable factorability. Secondly, the Kaiser-Meyer-Olkin measure of sampling adequacy was .683, above the commonly recommended value of .6, and Bartlett's test of sphericity was significant, $\chi^2(15) = 243.32$, $p < .001$. Therefore, we conducted a principal components analysis to identify and compute composite scores for the factors underlying the six measures. Initial eigen values indicated that the first two factors explained 58%, and 20% of the variance, respectively. The third, fourth, fifth, and sixth factors had eigen values below one, and

THE LESSER OF TWO EVILS

explained 13%, 5%, 2%, and 2% of the variance, respectively (for detailed results, see Table A2-1). The component matrix (Table A2-2) indicated that the first two components correspond to the three close others-related, and the three stranger-related questions, respectively, therefore we created two new variables (“attitudes towards close others” and “attitudes towards strangers”) for the two components by taking the arithmetic mean of the corresponding three measurements for each. These compound measurements indicated to what extent allocators cared about the beliefs and feelings of close others or strangers, respectively.

Table A2-1. Experiment 2 principal component analysis

Total Variance Explained									
Component	Initial Eigenvalues			Extraction Sums of Squared Loadings			Rotation Sums of Squared Loadings		
	Total	% of Variance	Cumulative %	Total	% of Variance	Cumulative %	Total	% of Variance	Cumulative %
1	3.583	59.715	59.715	3.583	59.715	59.715	2.563	42.718	42.718
2	1.121	18.676	78.391	1.121	18.676	78.391	2.14	35.673	78.391
3	0.707	11.783	90.174						
4	0.26	4.341	94.515						
5	0.24	4.004	98.519						
6	0.089	1.481	100						

Extraction Method: Principal Component Analysis.

Table A2-2. Experiment 2 component matrix

Component Matrix^a		
	Component	
	1	2
Fair to close others	0.756	0.542
Honest to close others	0.702	0.593
Feel good: close others	0.695	0.173
Fair to strangers	0.865	-0.379
Honest to strangers	0.829	-0.351
Feel good: strangers	0.775	-0.423

Extraction Method: Principal Component Analysis.

a. 2 components extracted.