



Case Report

Does observability amplify sensitivity to moral frames? Evaluating a reputation-based account of moral preferences

Valerio Capraro^{a,*}, Jillian J. Jordan^{b,1}, Ben M. Tappin^{c,1}^a Middlesex University London, United Kingdom^b Kellogg School of Management, Northwestern University, United States of America^c Sloan School of Management, Massachusetts Institute of Technology, United States of America

ARTICLE INFO

Keywords:

Moral preferences
 Moral frames
 Observability
 Trustworthiness
 Trust game
 Trade-off game

ABSTRACT

A growing body of work suggests that people are sensitive to moral framing in economic games involving prosociality, suggesting that people hold moral preferences for doing the “right thing”. What gives rise to these preferences? Here, we evaluate the explanatory power of a reputation-based account, which proposes that people respond to moral frames because they are motivated to look good in the eyes of others. Across four pre-registered experiments (total $N = 9601$), we investigated whether reputational incentives amplify sensitivity to framing effects. Studies 1–3 manipulated (i) whether moral or neutral framing was used to describe a Trade-Off Game (in which participants chose between prioritizing equality or efficiency) and (ii) whether Trade-Off Game choices were observable to a social partner in a subsequent Trust Game. These studies found that observability does not significantly amplify sensitivity to moral framing. Study 4 ruled out the alternative explanation that the observability manipulation from Studies 1–3 is too weak to influence behavior. In Study 4, the same observability manipulation *did* significantly amplify sensitivity to normative information (about what others see as moral in the Trade-Off Game). Together, these results suggest that moral frames may tap into moral preferences that are relatively deeply internalized, such that the power of moral frames is not strongly enhanced by making the morally-framed behavior observable to others.

1. Introduction

Humans are exceptionally prosocial. As one important illustration, many canonical experiments using economic games have shown that some people act prosocially even in one-shot anonymous interactions, when no direct or indirect benefits seem to be at play (see [Camerer \(2011\)](#) for a review). Understanding what drives such prosociality is essential for the success of human societies, given that we face global challenges such as resource depletion, climate change, social and economic inequalities, and pandemics ([Axelrod & Hamilton, 1981](#); [Hardin, 1968](#); [Nowak, 2006](#); [Rand and Nowak, 2013](#); [Perc et al., 2017](#); [Trivers, 1971](#); [Van Bavel et al., 2020](#)).

One classic explanation for human prosociality—and in particular, for prosociality in one-shot anonymous economic games—is that people have outcome-based social preferences ([Bolton & Ockenfels, 2000](#); [Charness & Rabin, 2002](#); [Engelmann & Strobel, 2004](#); [Fehr & Schmidt, 1999](#); [Levine, 1998](#)). This explanation poses that people care about the

monetary payoffs and outcomes of others, not just themselves (e.g., through preferences for others to receive resources, or for allocations of resources to be equal). Recent research, however, has suggested that beyond having outcome-based social preferences, people also have “moral preferences” for doing the right thing.

2. Moral frames and moral preferences

Over the last decade, a series of papers has revealed that people’s decisions in one-shot economic games depend on how the available choices are labelled, especially when the labels activate moral concerns. For example, [Krupka and Weber \(2013\)](#) found that the rate of altruism in the dictator game depends on how the game is framed: dictators tend to be more altruistic in the “take frame” compared to the “give frame”, despite the fact that the payoff outcomes for all parties are identical in both versions of the game. And this effect seems to be driven by moral considerations: dictators tend to rate “taking” from the recipient as more

* Corresponding author.

E-mail addresses: v.capraro@mdx.ac.uk (V. Capraro), jillianjordan@hbs.edu (J.J. Jordan), benmtappin@googlemail.com (B.M. Tappin).¹ These authors contributed equally to this work.

socially inappropriate than “not giving to” the recipient.

Furthermore, this type of effect has been documented across a wide range of economic game contexts. Framing effects have been documented for six dictator game frames by Capraro and Vanzo (2019). Eriksson, Strimling, Andersson, and Lindholm (2017) found that rejection rates in the ultimatum game depend on how the available actions are named, such that participants are more willing to reject an offer when doing so is described as “rejection of an offer” versus “reduction of a payoff”. And like Krupka and Weber (2013), Eriksson et al. (2017) also found evidence that this framing effect reflected moral considerations: “reducing the proposer’s payoff” was rated as morally worse than “rejecting the proposer’s offer”. Finally, Capraro and Rand (2018) and Tappin and Capraro (2018) found that minor changes in the framing of trade-off games pitting equity against efficiency can massively impact people’s decisions, at least when the payoffs are not too large (Huang, Lei, Xu, Yu, & Shi, 2019). And evidence suggests that, again, these framing effects can be explained by a change in the perception of what is the moral thing to do (Capraro & Rand, 2018). Together, these experiments robustly illustrate that beyond having outcome-based social preferences concerning the payoffs of others, people have preferences for doing the “right thing”.

These moral preferences provide a useful tool for intervening to promote prosociality. A growing body of work suggests that “moral suasion” (i.e., using frames that make the morality of an action salient) can function to increase socially desirable behavior. An earlier paper by Brañas-Garza (2007) found that telling dictators that “the other person relies on you” increases giving. Dal Bó and Dal Bó (2014) reported that reading about the Golden Rule encourages cooperation in the iterated prisoner’s dilemma. Capraro, Jagfeld, Klein, Mul, and van de Pol (2019) showed that asking participants “what do you personally think is the right thing to do?”, prior to their participation in a dictator game or a prisoner’s dilemma, increased prosociality in these games. Bilancini, Boncinelli, Capraro, Celadin, and Di Paolo (2020) found that moral suasion decreased ingroup favoritism on average (although it increased ingroup favoritism for a subset of participants).

The effect of moral suasion has also been observed in decisions that have consequences outside the laboratory. For example, Capraro et al. (2019) demonstrated that the aforementioned moral nudge, making salient that giving is the right thing to do, increases online crowd-sourcing of charitable donations to humanitarian organizations by 44%, and Bott, Cappelen, Sørensen, and Tungodden (2019) observed that letters with moral appeals, highlighting that paying taxes is fair towards other tax-payers and increase social benefits, decrease tax evasion among Norwegian tax-payers.

Thus, previous work has both demonstrated that people readily respond to moral frames—reflecting moral preferences for doing the right thing—and illustrated the potential for interventions to capitalize on moral preferences in order to encourage socially desirable behavior. However, much less is known about *why* people exhibit moral preferences and are sensitive to framing. What underlies the preference to do the “right” thing, and the resulting sensitivity to how choices are framed? Here, we consider this question, both from the perspective of the “ultimate” mechanisms through which moral preferences and framing effects may be adaptive, and the “proximate” mechanisms through which moral preferences and framing effects may be psychologically instantiated.

3. A reputation-based explanation for framing effects

One potential explanation for moral preferences and framing effects comes from considering the role of *reputation*. People are strongly motivated to be seen positively by others, and thus reputation is an important driver of prosocial behavior (Barclay, 2004; Boyd & Richerson, 1989; Emler, 1990; Milinski, Semmann, & Krambeck, 2002; Nowak & Sigmund, 2005). Supporting this proposal, evidence suggests that removing confidentiality increases public goods contributions

(Andreoni & Petrie, 2004; Rege & Telle, 2004) and charity donations (Bereczkei, Birkas, & Kerekes, 2007; Ariely, Bracha, & Meier, 2009; Lacetera & Macis, 2010). (See Bradley, Lawrence, and Ferguson (2018) for a recent meta-analysis.) Thus, one plausible reason for why people are sensitive to moral frames is that frames provide information about what will be seen as moral by *other* people; and, thus, responding to frames can confer reputational benefits. In other words, framing effects may draw on the reputation-based drive to appear moral in the eyes of others.

Yet, we know that framing effects can occur even in contexts where nobody is watching (e.g., one-shot anonymous economic game experiments). On its face, this observation is inconsistent with the proposal that framing effects draw on reputation motives: when nobody is watching, people are unlikely to be explicitly concerned with looking good in the eyes of others. However, evidence suggests that some people rely on the heuristic that reputation is typically at stake and use this heuristic even in one-shot game experiments where nobody is watching (Jordan & Rand, 2020). Such a heuristic might function to avoid the cognitive costs of computing whether an interaction is likely to be observed (Delton, Krasnow, Cosmides, & Tooby, 2011; Kahneman, 2011; Rand, Tomlin, Bear, Ludvig, & Cohen, 2017) or the social costs of performing such a computation—given that carefully calculating whether to cooperate or not can be perceived negatively by observers (Capraro & Kuilder, 2016; Critcher, Inbar, & Pizarro, 2013; Hoffman, Yoeli, & Nowak, 2015; Jordan, Hoffman, Nowak, & Rand, 2016). Furthermore, people might default towards behaving as if they are being watched as an “error management” strategy, designed to avoid the costs of mistaking a situation in which reputation is at stake for an anonymous situation (Delton et al., 2011).

Accordingly, it is plausible that framing effects *do* draw on reputation motives, but nonetheless occur in anonymous interactions because people are sensitive to reputation cues even when nobody is watching. Of course, a reputation-based explanation for framing effects might still predict that sensitivity to framing increases meaningfully in contexts where reputation is actually at stake, and people thus have explicit reputation motives. For example, Jordan and Rand (2020) found that the presence of reputational incentives can increase sensitivity to reputation cues among deliberative individuals (who are less likely to heuristically assume that reputation is at stake even in anonymous interactions). Thus, to the extent that framing effects draw on reputation concerns, we might expect them to increase when explicit reputational incentives are introduced.

Here, we evaluate this prediction by investigating the extent to which sensitivity to framing is itself sensitive to whether or not reputation is at stake. This inquiry has important practical implications: if sensitivity to framing is amplified substantially by the presence of explicit reputational incentives, it suggests that there is substantial value in targeting framing-based interventions at contexts where reputation is at stake. Furthermore, investigating the effect of reputational incentives on sensitivity to framing can shed light on the mechanisms underlying framing effects. Specifically, such an investigation can shed light on the extent to which framing effects draw on psychological processes that are context-sensitive (and thus increase substantially when explicit reputational incentives are introduced) versus deeply internalized (and thus relatively insensitive to whether we are being watched).

4. Our contribution

In this paper, we investigate the influence of reputational incentives on sensitivity to moral frames in the context of the Trade-Off Game (TOG), an economic decision problem that forces decision-makers to decide between competing values of equality and efficiency. We chose to investigate decisions in the TOG because there is a robust body of evidence that, in experiments where reputation is *not* at stake, moral framing has a substantial influence on TOG decisions (Capraro & Rand, 2018; Capraro, Rodriguez-Lara, & Ruiz-Martos, 2020; Huang et al.,

2019; Tappin & Capraro, 2018)—and, crucially, on moral judgments (Capraro & Rand, 2018), such that the vast majority of participants report that the moral thing to do in the TOG is whichever action is targeted by the moral frame.

We thus conducted four pre-registered experiments in which participants (i) decided how to behave in the TOG, and then (ii) participated in an economic Trust Game with another player. As described further below, participants participated in the Trust Game in the role of the “Trustee”, and thus faced incentives to appear trustworthy in the eyes of the other player.

Studies 1–3 employed a two-by-two, between-subjects design in which we manipulated (i) the framing used to describe participants’ TOG decisions (neutral frame vs. moral frame) and (ii) whether or not participants’ TOG decisions were observable to the other player in the Trust Game (private condition vs. public condition). This design allows us to investigate the extent to which activating reputation concerns (by making TOG decisions observable) increases sensitivity to moral framing. Validating this reputation manipulation, previous work has demonstrated that the desire to be evaluated positively in the Trust Game provides a salient reputation motivation that can shape behavior in the moral domain (e.g., Jordan, Hoffman, Bloom, & Rand, 2016; Jordan, Hoffman, Nowak, & Rand, 2016; Jordan & Rand, 2020).

Yet, to preview our results, Studies 1–3 reveal that making TOG decisions observable does *not* significantly amplify sensitivity to moral framing. Thus, in Study 4, we sought to provide evidence against the alternative explanation that the reputation manipulation from Studies 1–3 is simply too weak to meaningfully influence behavior in the context of the TOG. To this end, in Study 4 we replaced our moral framing manipulation with a manipulation of normative information (about what others see as moral in the Trade-Off Game). Study 4 shows that the reputation manipulation from Studies 1–3 *does* amplify sensitivity to this normative information, suggesting that our reputation manipulation is *not* simply too weak to shape behavior in the TOG.

In the Trust Game, we also measured the extent to which focal participants actually behaved in a trustworthy way towards the other player. Thus, our design also allows us to ask an additional question: to what extent do individuals who make morally framed choices in the TOG behave in a more trustworthy way in the Trust Game? Previous work has found that individuals who make morally framed choices in the TOG tend to behave more prosocially in subsequent one-shot anonymous economic games like the Dictator Game and Prisoner’s Dilemma (Capraro & Rand, 2018; Tappin & Capraro, 2018). These results further support the proposal that sensitivity to framing reflects the prosocial preference to do the right thing. And they lend credulity to the reputation-based hypothesis we explore in our primary analyses: Insofar as morally-framed choices are a meaningful signal of prosociality, it makes sense that the desire to appear prosocial in the eyes of others might serve to amplify sensitivity to moral frames. Thus, as a secondary question of interest in our paper, we sought to replicate the association between morally-framed choices and prosociality in the context of our paradigm. To this end, we investigated whether selecting the morally-framed choice in the TOG would serve to signal trustworthiness in the Trust Game.

4.1. Studies 1–2

We begin by presenting the results of Studies 1–2. Owing to their close similarity, we present the results of these two studies together. The study designs, hypotheses, and analysis plans were pre-registered. The protocols are online at <https://aspredicted.org/yz4gk.pdf> (Study 1) and <https://aspredicted.org/zz5my.pdf> (Study 2).

5. Method

5.1. Samples

Participants were recruited via Amazon’s Mechanical Turk (AMT, Paolacci, Chandler, & Ipeirotis, 2010; Horton, Rand, & Zeckhauser, 2011; Paolacci & Chandler, 2014; Arechar, Gächter, & Molleman, 2018). The sample sizes and payment information are reported below (subheading “Study 1 and Study 2”).

5.2. Procedure and design

Participants were randomly assigned to one of four conditions, corresponding to a two-by-two, between-subjects design in which we manipulated Choice Frame (Neutral, Moral) and Choice Observability (Public, Private). We describe each of these manipulations in the context of our experimental procedure below.

5.2.1. Trade-off game

Participants began the study by learning that they would be playing a game called the “Trade-Off Game” (TOG). The TOG instructions described to participants that they would be playing as “Player A”, and that they would be matched with two other players, “Player B” and “Player C”. As Player A, participants were asked to choose between one of two options: Option 1 (referred to by us in this paper as the equitable option) and Option 2 (referred to by us in this paper as the efficient option). Option 1 involved providing 13 cents to each of the three players (Player’s A, B, and C). In contrast, Option 2 involved providing 13 cents to Player A, 23 cents to Player B and 13 cents to Player C. Thus, Option 2 provided a larger total payoff for the group (and was thus desirable from an efficiency standpoint), but also conferred an unequal distribution of money to group members (and was thus undesirable from an equality standpoint).

In the Neutral Frame condition, we labelled the options simply as “Option 1” and “Option 2”. In contrast, in the Moral Frame condition, we labelled Option 1 (i.e., the option that values equality) as the “fair” option. Supporting this design, previous research has established that, in this Moral Frame condition, the vast majority of participants indeed think that selecting Option 1 is the morally right thing to do (Capraro & Rand, 2018). We also note that we could have contrasted our Moral Frame condition not with the Neutral Frame condition, but rather with a “reverse” Moral Frame condition (that framed Option 2 as the morally right thing to do). Such a design would have created a relatively larger contrast between our two framing conditions. Thus, we opted for our design so that (i) our framing manipulation would have a relatively smaller effect on behavior in private and (ii) there would thus be more “room” for the public condition to amplify sensitivity to framing (i.e., there would be less risk of a ceiling effect).

After reading the TOG instructions, all participants were asked two comprehension questions to assess their understanding of the task.

5.2.2. Trust game

After completing the TOG comprehension questions, participants were informed that they would be playing another game: The Trust Game (TG). All participants were informed they would be playing the TG in the role of the “Receiver” (i.e., the Trustee), and that a new player (who was *not* involved in the TOG) would be playing the TG in the role of the “Sender” (i.e., the Trustor). Participants were then told that the Sender would start with 30 cents and choose how many cents, if any, to send to them. Each cent sent would be tripled before being given to the participant, who would then get to decide how much, if anything, to return to the Sender.

In the *Public* condition, participants were told that the Sender would learn about the existence of the TOG and find out about their TOG decision. In contrast, in the *Private* condition, participants were told that the Sender would *not* learn about the existence of the TOG or find out

about their TOG decision. Thus, in the Public condition, participants knew that their Trust Game partner would observe their TOG decision before deciding how much money to entrust them with, whereas participants in the Private condition knew that their TOG decision would not be observed.

Following the TG instructions, participants were asked four comprehension questions to assess their understanding of the task.

5.2.3. Game decisions

After completing the TG comprehension questions, all participants made their game decisions. First, they decided whether to make the equitable or efficient choice in the TOG, and second, they decided how much to return (as a proportion of whatever amount they would be sent) to their partner in the TG. In the TOG decision screen, participants were reminded whether their choice would (Public condition) or would not (Private condition) be revealed to the Sender. The TG decision screen was the same for all the participants. Verbatim instructions for both the TOG and TG are reported in the Supplementary Information. Finally, after making their decisions, participants completed a demographic questionnaire.

5.2.4. Study 1 and Study 2

As described, we collected data in two pre-registered studies. In Study 1, the participation payment was 50 cents and we pre-registered a target sample size of $N = 800$ participants. In Study 2, we pre-registered a much larger target sample size of $N = 3000$. The procedure for Study 2 was identical to the procedure for Study 1, with the exception that (i) in Study 2 the participation payment was 30 cents (instead of 50 cents) and we correspondingly shortened the demographic questionnaire to reduce the length of the survey, and (ii) in Study 2 the stakes in the TOG were lower. Specifically, the stakes in the TOG were [5,5,5] vs. [5,10,5] cents (instead of [13,13,13] vs. [13,23,13]), and, in the TG, the Sender started with 10 cents (instead of 30 cents).

5.2.5. Disclosure statement

In these studies, we report all measures, manipulations and exclusions.

5.2.6. Sensitivity power analysis

We conducted sensitivity power analyses for our primary test (Hypothesis 1) via simulations. The script to reproduce our simulations is available online on <https://osf.io/45cey/>. For all simulations, we assume a moral framing effect in the private condition equal to an increase of 15 percentage points and a baseline of 50% (i.e., neutral-private condition proportion = 0.50; moral-private condition proportion = 0.65). Accordingly, Study 1 ($N = 800$) had 80% power to detect an interaction between the moral frame and public condition assuming that the effect of the moral frame in public was approximately double the size of that observed in private: neutral-public condition proportion = 0.50; moral-public condition proportion = 0.82. Owing to its larger sample size ($N = 2810$), Study 2 had 80% power to detect a smaller increase in the moral frame effect in public: neutral-public condition proportion = 0.50, moral-public condition proportion = 0.75. The pooled data from Studies 1 and 2 had approximately 80% power to detect a slightly smaller effect: neutral-public condition proportion = 0.50, moral-public condition proportion = 0.74.

5.2.7. Ethics statement

For all the sessions, relevant ethical guidelines were followed. This research was approved by the MIT IRB, COUHES Protocol #: 18066392996A004.

6. Results

Data and analysis code to reproduce all results is available online at <https://osf.io/45cey/>.

For all analyses, we report results among all participants (preregistered primary analyses) and among participants who correctly answered comprehension questions about both economic games (preregistered exploratory analyses: briefly described towards the end of this section, but reported in full in the SI).

We preregistered two data exclusion criteria: We exclude responses from participants who did not complete all of the primary measures prior to the demographics (because these participants by definition could not be included in the preregistered analyses). We also exclude those who had duplicate IP addresses or AMT worker IDs, retaining the first response only (determined by the start date). Our final sample size is thus $N = 3610$ (Study 1 $N = 800$, Study 2 $N = 2810$). Our preregistered analyses address two key questions.

6.1. Analysis 1: Does observability increase sensitivity to moral framing in the TOG?

Our first and primary analysis asks whether people are more sensitive to the moral framing manipulation when their choice is observable (versus unobservable) by their partner in the subsequent Trust Game. Thus, we test the two-way interaction between framing (neutral, moral) and observability (private, public) on TOG choices. The relevant raw data are displayed in Fig. 1A, partitioned by study.

We analyze these data using binomial logistic regression. The results for each study are reported in Table 1 (Analysis 1). Replicating past work and illustrating participants' preferences to behave morally, we find consistent evidence across studies for a main effect of the moral frame: Participants are more likely to choose the equitable (vs. efficient) option in the TOG when that choice is framed as the "fair" option than when there is no such frame. This effect can be seen in Fig. 1A by comparing results in the moral versus neutral frame conditions. In Table 2 (Analysis 1), we pool the data across the studies—including a dummy variable to indicate study-specific differences in the intercept—to obtain a more precise estimate of the moral framing estimates.

Next, we turn to our key question: Does observability increase sensitivity to moral framing? We do not find strong evidence that it does. That is, our results imply that participants are sensitive to moral framing both in *private* (where reputational incentives are absent, and participants are unlikely to be explicitly motivated by reputation concerns) and in *public* (where reputational incentives are present, and participants may be explicitly motivated by reputation concerns). And while the observed framing effect is directionally larger in public than in private (as illustrated in Fig. 1A), we do not find evidence for a significant amplification in the presence of reputational incentives. Table 1 (Analysis 1) reveals that we do not find a significant interaction between our moral frame and observability manipulations in either of our studies individually, and Table 2 (Analysis 2) reveals that pooling data across studies to maximize precision gives the same result.

Thus, the results fail to provide strong evidence that observability increases sensitivity to moral framing, and thus suggest that moral framing may be similarly effective in the absence versus presence of reputational incentives.

6.2. Analysis 2: Do Morally-Framed Decisions Signal Trustworthiness?

Our second analysis asks whether morally-framed decisions signal trustworthiness, and is broken into two parts. To address this question, we first ask in Analysis 2a whether participants who choose the morally-framed option (i.e., the equitable choice) in the TOG behave in a more trustworthy manner in the subsequent Trust Game. Then, in Analysis 2b, we turn to asking whether any such association depends on the presence of the moral frame, or is instead explained by a general association between equitable choices and trustworthiness (that holds irrespective of framing). The relevant raw data are shown in Fig. 1B.

In these analyses, we investigate the association between morally-framed choices and trustworthiness in the *Private* condition of our

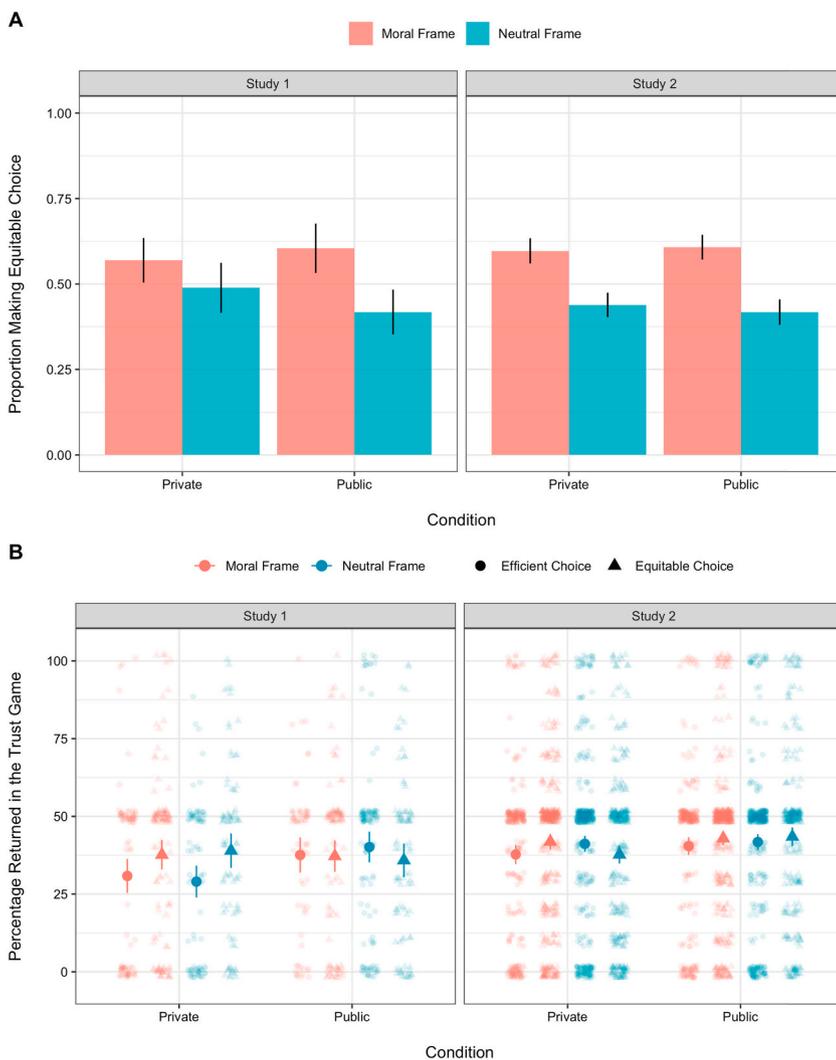


Fig. 1. Data from Studies 1 and 2. (A) Proportion of participants choosing to split the money equitably in the Trade-Off Game as a function of observability (Private, Public) and choice frame (Moral, Neutral). (B) Percentage of money returned in the Trust Game as a function of observability (Private, Public), choice frame (Moral, Neutral), and choice made in the Trade-Off Game (Efficient, Equitable). Large solid points are the means in each cell. Smaller faded points are the raw data with slight jitter to aid visibility. (A, B) Error bars are 95% CI. $N = 800$ in Study 1; $N = 2810$ in Study 2.

paradigm, in which reputational incentives were absent and thus sensitivity to moral framing can only reflect “pure” preferences for doing the right thing. This approach has the advantage of allowing us to straightforwardly compare our results to previous work investigating the association between morally-framed choices and prosociality, which has employed one-shot anonymous games that are analogous to our Private condition (Capraro & Rand, 2018; Tappin & Capraro, 2018). These analyses were thus preregistered as primary analyses because we had clear theoretical predictions. Because we did not have clear theoretical predictions about how the Public condition might affect the association between morally-framed choices and trustworthiness, we preregistered analyses on the Public condition as exploratory, and we prioritize brevity by not additionally reporting those analyses here. However, we note that interested readers can get a visual sense of the relevant associations from Fig. 1B, and we report formal results from these exploratory analyses on the Public condition specifically in the SI.

6.2.1. Analysis 2a

We first examine the trustworthiness of participants in the Private condition who received the moral frame in the TOG. This functions as a basic test of whether participants who choose the equitable option behaved in a more trustworthy manner than participants who choose the efficient option, when the equitable option is framed as morally good and decisions are private. Understood with respect to Fig. 1B, this test compares the first two cells of the Private condition. We analyze these data using OLS regression. The results for each study are reported in

Table 1 (Analysis 2a), and the pooled estimate across both studies is reported in Table 2 (Analysis 2a).

We find some evidence that participants who chose the equitable option returned more in the subsequent Trust Game than participants who chose the efficient option: the estimates in both studies are positive, and the p -values on the estimates are 0.067 (Study 1) and 0.041 (Study 2). Thus, we find some evidence that equitable choices predict trustworthiness—when such choices are framed as morally good and all decisions are private. The effect size is modest, but not tiny: as per Table 2 (Analysis 2a), participants who chose equitably in the TOG returned on average ~ 4.8 percentage points more in the Trust Game than those who chose the efficient option. Compared against the return rate of participants who chose the efficient option (~ 32 percentage points, as per the Intercept in Table 2 Analysis 2a) the increase in return rate that we observe among participants who chose equitably is approximately +15% (i.e., $4.8/32$).

6.2.2. Analysis 2b

We next proceed to ask: Is the association identified in Analysis 2a observed only when the equitable choice is framed as morally good, or does it instead reflect a general association between equitable choices and trustworthiness (that holds irrespective of framing)?

To address this question, we investigate the two-way interaction between framing (neutral, moral) and TOG choice (equitable, efficient) on trustworthiness in the TG. Understood with respect to Fig. 1B, this test compares the difference between the first two cells of the Private

Table 1
Results from Analyses 1 and 2 in each of Studies 1 and 2.

Analysis	Study	Term	Estimate	95 LL	95 UL	p-value	
1	1	(Intercept)	0.96	0.71	1.28	0.766	
		Moral Frame	1.38	0.93	2.05	0.107	
		Public Condition	0.75	0.51	1.12	0.158	
	2	Moral x Public	1.54	0.87	2.70	0.135	
		(Intercept)	0.78	0.68	0.90	0.001	
		Moral Frame	1.89	1.53	2.34	<	
2a	2	Public Condition	0.92	0.74	1.13	0.423	
		Moral x Public	1.14	0.84	1.54	0.391	
		(Intercept)	30.83	25.36	36.31	<	
2a	1	Equitable Choice	6.80	-0.45	14.06	0.067	
		(Intercept)	37.75	34.71	40.78	<	
		Equitable Choice	4.10	0.17	8.03	0.041	
2b	1	(Intercept)	29.02	23.57	34.48	<	
		Equitable Choice	9.96	2.15	17.76	0.013	
		Moral Frame	1.81	-5.82	9.45	0.642	
	2	Equitable x Moral	-3.15	-13.69	7.38	0.558	
		(Intercept)	41.17	38.65	43.69	<	
		Equitable Choice	-3.47	-7.27	0.34	0.074	
	2b	2	Moral Frame	-3.43	-7.43	0.57	0.093
			Equitable x Moral	7.56	2.03	13.1	0.007
			Moral				

Note. Analysis 1 model estimates and CIs are odds ratios. Analysis 2a and 2b are OLS estimates. LL = Lower Limit. UL = Upper Limit.

Table 2
Results from Analyses 1 and 2 pooling the data across Studies 1 and 2.

Analysis	Term	Estimate	95 LL	95 UL	p-value
1	(Intercept)	0.82	0.69	0.99	0.037
	Moral Frame	1.77	1.47	2.13	< 0.001
	Public Condition	0.88	0.73	1.06	0.184
	Study 2 Dummy	0.99	0.84	1.16	0.854
	Moral x Public	1.22	0.93	1.59	0.145
2a	(Intercept)	31.99	28.04	35.94	< 0.001
	Equitable Choice	4.77	1.31	8.23	0.007
	Study 2 Dummy	5.35	1.40	9.30	0.008
2b	(Intercept)	34.45	31.13	37.77	< 0.001
	Equitable Choice	-0.83	-4.25	2.59	0.635
	Moral Frame	-2.59	-6.13	0.95	0.152
	Study 2 Dummy	5.52	2.59	8.45	< 0.001
	Equitable x Moral	5.60	0.70	10.50	0.025

Note. Analysis 1 model estimates and CIs are odds ratios. Analysis 2a and 2b are OLS estimates. LL = Lower Limit. UL = Upper Limit.

condition with the difference between the latter two cells of the Private condition. We analyze these data using OLS regression. The results for each study are reported in Table 1 (Analysis 2b), and the pooled estimate across both studies is reported in Table 2 (Analysis 2b).

We find some evidence that participants who chose the equitable (versus efficient) option returned more in the subsequent Trust Game *only* when the equitable option was framed as morally good. Specifically, while there is no support for this pattern (even directionally) in Study 1, the evidence does support it in Study 2, and in the pooled estimate across both studies (as indicated by the significant positive interaction in Table 2, Analysis 2b). Thus, we find some evidence that, when decisions are private, equitable choices specifically predict trustworthiness when they are framed as morally good.

Together, the results of Analysis 2 are thus broadly supportive of the proposal that, in a context where reputational incentives are absent—and thus sensitivity to moral framing can only reflect “pure” moral preferences—morally-framed choices possess some value as a signal of trustworthiness.

6.3. Exploratory analyses among perfect comprehenders

Finally, as mentioned, we also conducted the above analyses (i.e., 1, 2a, and 2b) on a restricted sample of participants: those who correctly answered all comprehension questions for the TOG and TG (with the exception that one question was not considered when performing this restriction²). We refer to this set of participants as “perfect comprehenders”. Their results are reported in figures and tables in the SI. The main result of note is that, in Study 2, Analysis 1, we *do* observe a statistically significant interaction between moral frame and public condition ($p = .048$), reflecting that the moral framing effect on TOG choice is larger in public than private. This result suggests that, among people who fully understand the choice context, observability may in fact increase sensitivity to moral framing in the TOG. However, because this was an exploratory finding, it may alternatively reflect a type I error. Thus, we conducted a third study: a preregistered direct replication of Analysis 1, focused on perfect comprehenders only.

6.3.1. Study 3

This study was pre-registered at: <https://aspredicted.org/8ge6a.pdf>. The design was identical to Study 2, including the targeted sample size ($N = 3000$ participants recruited from AMT). Participants were unable to take part in Study 3 if they had completed either of Studies 1 or 2. Besides the addition of restricting to perfect comprehenders only, Study 3 had the same preregistered exclusion criteria as Studies 1 and 2. The relevant raw data are shown in Fig. 2 under heading “perfect comprehenders” (we also show the data for all participants). As illustrated by the figure, the effect of moral framing is comparable in private versus public. Confirming this visual impression, logistic regression reveals that perfect comprehenders in Study 3 show no statistically significant interaction between observability and framing, $OR = 1.07 [0.70, 1.62]$ $p = .763$. (We also find no significant interaction when analyzing all participants, $p = .610$, see SI for full tables). (A separate finding of note is that Study 3 does not replicate our findings from Studies 1–2 regarding Analysis 2a; the point-estimate is of the same sign, but $p > .05$, see SI for details.)

Thus, the results of Study 3 suggest that the statistically significant interaction that we observed among perfect comprehenders in Study 2 was likely a false positive. More generally, the results of Study 3 align with the results of Studies 1 and 2: there is a lack of evidence that observability increases sensitivity to moral framing in the TOG. Further bolstering this conclusion, pooling the data from all participants in Studies 1–3 ($N = 6601$) reveals that the moral framing effect is approximately 91% as large in the Private condition as in the Public condition, and that there is no statistically significant interaction between observability and framing ($p = .444$).

6.3.2. Study 4

Together, Studies 1–3 suggest that the influence of moral framing on TOG decisions is relatively insensitive to observability. However, an alternative explanation is that the observability manipulation used in Studies 1–3 is simply too weak to meaningfully shape behavior. We

² By necessity, the comprehension questions come after the instructions for the TOG and TG, and thus also come after treatment assignment. For one of our TG questions, the correct answer is different in the public vs. private condition. And for this question, the pass rate differed significantly (at $p < .05$) between conditions. Thus, to minimize potential post-treatment bias, when defining “perfect comprehension”, we do not use this comprehension question.

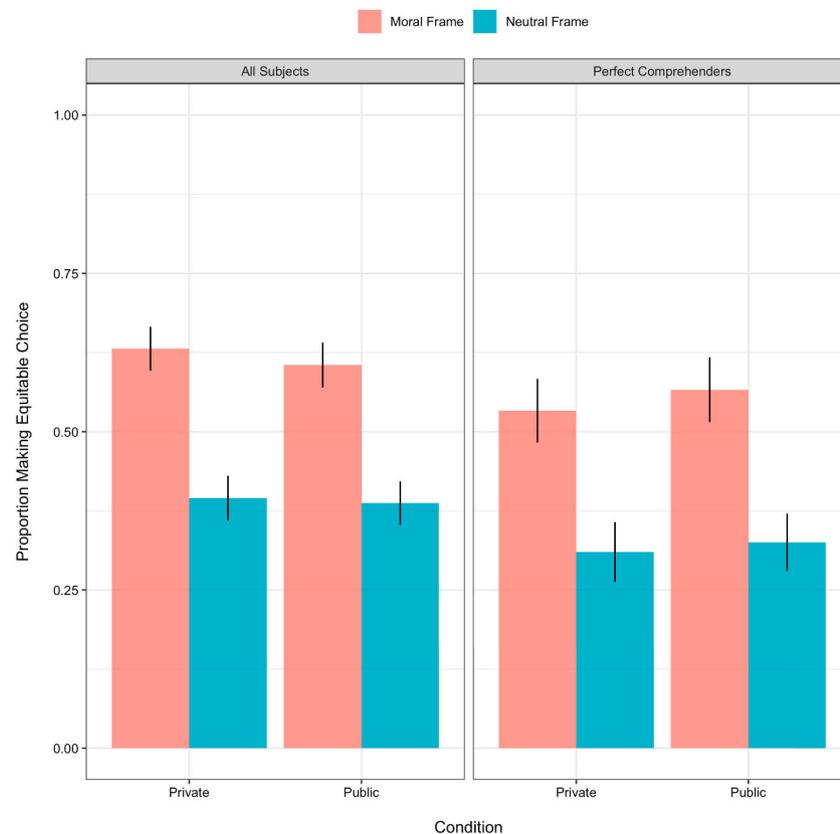


Fig. 2. Results of Study 3: a direct replication of Study 2. $N = 2991$ all subjects; $N = 1525$ perfect comprehenders.

empirically test this possibility in a fourth study, pre-registered at: <https://aspredicted.org/y32m3.pdf>.³ In Study 4, we use the same observability manipulation from Studies 1–3. However, instead of asking whether it amplifies sensitivity to *moral framing*, we ask whether it amplifies sensitivity to *normative information* about what others see as moral in the Trade-Off Game. We reasoned that if our observability manipulation is psychologically impactful (such that subjects in the public condition are in fact motivated to be evaluated by the Trust Game sender), it should heighten motivation to conform to such normative information.

Thus, the design of Study 4 was identical to the design of Studies 1–3, with one key difference: instead of manipulating moral framing (by varying whether the equitable TOG choice was labelled as “fair”), we manipulated normative information. Specifically, we varied whether participants were told about the TOG decision of ten participants from a previous survey. In this group of ten, nine (90%) said that the equitable option was morally right.⁴ Thus, Study 4 still manipulated whether subjects were encouraged to choose the equitable option—but via normative information, rather than a moral frame. We investigated whether our observability manipulation would amplify sensitivity to

this normative information. We refer to the SI for experimental instructions.

The target sample size of Study 4 was $N = 3000$, and the preregistered exclusion criteria were the same as Studies 1–3. Participants who completed Study 3 were unable to take part in Study 4. The raw data are displayed in Fig. 3. As illustrated by the figure, and in contrast to the results from Studies 1–3, the effect of normative information is meaningfully larger in public than in private—approximately twice as large. Indeed, logistic regression reveals a statistically significant interaction between observability among all participants, $OR = 1.40 [1.05, 1.87]$ $p = .024$ (our primary pre-registered analysis). We note, however, that the interaction does not reach significance (at $p < .05$) in an exploratory analysis of perfect comprehenders, $p = .139$ (but the point estimate is the same, $OR = 1.39$); see SI for full tables.

These results negate the proposal that our observability manipulation is simply too weak to influence behavior in the TOG. And in this way, they suggest that Studies 1–3 do in fact provide evidence that moral framing effects are relatively insensitive to the influence of reputation.

7. Discussion

A growing body of research has illustrated that people are sensitive to framing when making moral decisions. This research suggests that beyond having outcome-based social preferences (e.g., for others to receive resources, or for resource allocations to be equal), people also have preferences for doing the “right” thing. In this work, we have investigated the extent to which reputational incentives serve to amplify sensitivity to moral framing effects, and thus the extent to which the moral preferences underlying moral framing effects are sensitive to whether reputation is at stake. And we have provided evidence that moral framing effects are relatively insensitive to reputation. In Studies 1–3, participants responded to frames even when their decisions were

³ Unfortunately, the pre-registration for Study 4 was approved after collecting the data due to a miscommunication: one of the authors did not click the approval button prior to data collection, despite approving of the registration verbally prior to data collection with the other two authors (who both did click the approval button prior to data collection). We are happy to provide on request time-stamped email correspondences confirming the verbal agreement and pre-registration submission that occurred prior to data collection.

⁴ Capraro and Rand (2018), Study 4, asked subjects to report what they think is the morally right thing to do in the trade-off game in several frames, including the “fair” frame used in the current paper. 77% of the subjects said that being equitable was morally right. This means that we can draw 10 subjects such that 9 of them said that being equitable was morally right.

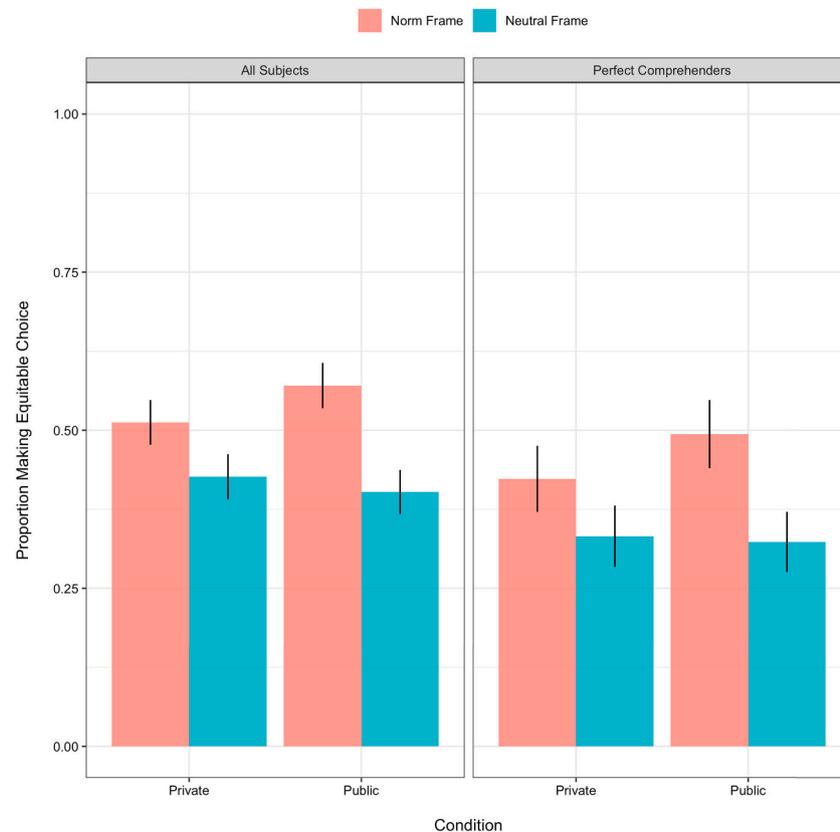


Fig. 3. Results of Study 4, in which our manipulation of moral framing was replaced with a manipulation of normative information.

completely private, and we did not find reliable evidence that reputation increases sensitivity to framing. Furthermore, Study 4 suggests that our reputation manipulation was *not* simply too weak to shape behavior: the very same manipulation amplifies sensitivity to *normative* information.

We note that our results are not inconsistent with the possibility that reputational incentives have some limited power to increase sensitivity to moral frames. Directionally, we did find larger framing effects when reputation was at stake; pooling the data from Studies 1–3, the moral framing effect was 91% as large in the Private condition as in the Public condition. Yet, even with an N of 6601 (pooling across Studies 1–3), we were not able to detect statistically significant amplification. We thus see our results as providing strong evidence that moral framing effects are relatively insensitive to the presence of reputational incentives, and that reputational incentives do not strongly amplify sensitivity to moral frames.

Our studies do, however, lend further support to the proposal that morally-framed choices serve as a reliable signal of trustworthiness. In our secondary analyses, we found some evidence that equitable decisions in the Trade-off Game predict returning in the Trust Game—but only when such choices are framed as morally good. This is consistent with previous evidence that morally-framed decisions predict prosocial behavior (Capraro & Rand, 2018; Tappin & Capraro, 2018), and serve to further bolster the proposal that sensitivity to moral framing reflects the moral preference to do the right thing. They also lend credulity to the reputation-based hypothesis we investigated in our primary analyses by suggesting that sensitivity to moral framing may be an effective reputational strategy (insofar as observers are aware that morally-framed choices signal prosociality). Thus, it is all the more notable that sensitivity to framing is relatively insensitive to the presence of reputational incentives.

Our results have implications for interventions that draw on moral framing effects to encourage socially desirable behavior. They suggest that such interventions can be successful even when behavior is not

observable to others and thus reputation is not at stake—and in fact, that the efficacy of moral framing effects is not strongly enhanced by making behavior observable. Thus, our results suggest that targeting contexts where reputation is at stake is not an especially important priority for individuals seeking to maximize the impact of interventions based on moral framing. This conclusion provides an optimistic view of the potential of such interventions, given that there may be many contexts in which it is difficult to make behavior observable but yet possible to frame a decision in a way that encourages prosociality—for example, when crowdsourcing donations anonymously (or nearly anonymously) on the Internet. Future research should investigate the power of moral framing to promote prosocial behavior in anonymous contexts outside of the laboratory.

Interestingly, Study 4 shows that normative information likewise has an effect on private choices (see also Human & Capraro, 2020, for a replication), but, importantly, the influence of normative information *does* become significantly larger when reputation is at stake. Thus, in contrast to moral framing, interventions that use normative information to encourage socially desirable behavior may be particularly impactful in public.

Our results also have important implications for the psychology underlying moral framing effects. Specifically, they suggest that moral framing effects draw on preferences that are relatively deeply internalized (and thus are relatively insensitive to the presence of reputational incentives). This observation could reflect that reputation systems do *not* explain why people are sensitive to moral framing. In other words, our results could reflect that people react to moral framing not because they see frames as conveying information about what others will see as moral, but rather because they see frames as conveying information about what actually is moral—and frames tap into a general and deeply internalized desire to behave morally (and/or think of the self as behaving morally) (Ariely, Bracha, & Meier, 2009; Mazar, Amir, & Ariely, 2008). Such a desire would serve to encourage cooperative and

moral behavior, and thus could be supported by ultimate mechanisms beyond reputation systems that can give rise to such behavior (e.g., direct reciprocity, kin selection, institutions, and cultural group selection) (Boyd & Richerson, 2009; Henrich, 2006; Nowak, 2006; Trivers, 1971).

Alternatively, however, our results could be compatible with the proposal that reputation systems explain why people are sensitive to moral framing. Specifically, our results could reflect that people do react to moral framing because moral frames convey information about what others will see as moral—and consequently do tap into a more specific desire to appear moral in the eyes of others that is ultimately supported by reputation systems. However, insofar as reputation is the primary driver of moral framing effects, our results suggest that the relevant reputation motives are strongly activated even in anonymous experiments, such that introducing making behavior observable does not substantially amplify sensitivity to framing. This possibility is consistent with the proposal that even when nobody is watching and people are not explicitly concerned with their reputations, they may nonetheless be implicitly motivated to appear virtuous—reflecting the use of heuristics (specifically, the heuristic that reputation is typically at stake) (Jordan & Rand, 2020) and/or error management strategies (specifically, strategies designed to avoid the costs of mistaking a situation in which reputation is at stake for an anonymous situation) (Delton et al., 2011).

Regardless, our results suggest that moral framing effects function by drawing on preferences that are not highly sensitive to the presence of reputational incentives. And, consequently, they suggest that sensitivity to moral framing effects—and the underlying preference to behave morally—may be considered a core and deeply internalized feature of our moral psychology. This conclusion implies that we might expect moral framing effects to be relatively robust and invariant across contexts, a prediction that should be probed further in future research. In particular, an interesting question is whether reputation might be more effective at amplifying the power of framing in contexts where baseline sensitivity to framing is relatively low (e.g., because the morally framed decision is particularly costly or difficult).

Acknowledgement

This paper has been funded from the John Templeton Foundation grant #61061.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jesp.2021.104103>.

References

- Andreoni, J., & Petrie, R. (2004). Public goods experiments without confidentiality: A glimpse into fund-raising. *Journal of Public Economics*, 88(7–8), 1605–1623.
- Arechar, A. A., Gächter, S., & Molleman, L. (2018). Conducting interactive experiments online. *Experimental Economics*, 21, 99–131.
- Ariely, D., Bracha, A., & Meier, S. (2009). Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *American Economic Review*, 99, 544–555.
- Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211, 1390–1396.
- Barclay, P. (2004). Trustworthiness and competitive altruism can also solve the “tragedy of the commons”. *Evolution and Human Behavior*, 25(4), 209–220.
- Bereczkei, T., Birkas, B., & Kerekes, Z. (2007). Public charity offer as a proximate factor of evolved reputation-building strategy: An experimental analysis of a real-life situation. *Evolution and Human Behavior*, 28, 277–284.
- Bilancini, E., Boncinelli, L., Capraro, V., Celadin, T., & Di Paolo, R. (2020). “Do the right thing” for whom? An experiment on ingroup favoritism, group assorting and moral suasion. *Judgment and Decision making*, 15, 182–192.
- Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, 90, 166–193.
- Bott, K. M., Cappelen, A. W., Sørensen, E.Ø., & Tungodden, B. (2019). You’ve got mail: A randomized field experiment on tax evasion. *Management Science*.
- Boyd, R., & Richerson, P. J. (1989). The evolution of indirect reciprocity. *Social Networks*, 11(3), 213–236.
- Boyd, R., & Richerson, P. J. (2009). Culture and the evolution of human cooperation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1533), 3281–3288.
- Bradley, A., Lawrence, C., & Ferguson, E. (2018). Does observability affect prosociality? *Proceedings of the Royal Society B: Biological Sciences*, 285, 20180116.
- Brañas-Garza, P. (2007). Promoting helping behavior with framing in dictator games. *Journal of Economic Psychology*, 28, 477–486.
- Camerer, C. F. (2011). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
- Capraro, V., Jagfeld, G., Klein, R., Mul, M., & van de Pol, I. (2019). Increasing altruistic and cooperative behaviour with simple moral nudges. *Scientific Reports*, 9, 11880.
- Capraro, V., & Kuilder, J. (2016). To know or not to know? Looking at payoffs signals selfish behavior, but it does not actually mean so. *Journal of Behavioral and Experimental Economics*, 65, 79–84.
- Capraro, V., & Rand, D. G. (2018). Do the right thing: Experimental evidence that preferences for moral behaviour, rather than equity or efficiency per se, drive human prosociality. *Judgment and Decision making*, 13, 99–111.
- Capraro, V., Rodriguez-Lara, I., & Ruiz-Martos, M. J. (2020). Preferences for efficiency, rather than preferences for morality, drive cooperation in the one-shot stag-hunt game. *Journal of Behavioral and Experimental Economics*.
- Capraro, V., & Vanzo, A. (2019). The power of moral words: Loaded language generates framing effects in the extreme dictator game. *Judgment and Decision making*, 14, 309–317.
- Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117, 817–869.
- Critcher, C. R., Inbar, Y., & Pizarro, D. A. (2013). How quick decisions illuminate moral character. *Social Psychological and Personality Science*, 4, 308–315.
- Dal Bó, E., & Dal Bó, P. (2014). “Do the right thing”: The effects of moral suasion on cooperation. *Journal of Public Economics*, 117, 28–38.
- Delton, A. W., Krasnow, M. M., Cosmides, L., & Tooby, J. (2011). Evolution of direct reciprocity under uncertainty can explain human generosity in one-shot encounters. *Proceedings of the National Academy of Sciences*, 108, 13335–13340.
- Emler, N. (1990). A social psychology of reputation. *European Review of Social Psychology*, 1(1), 171–193.
- Engelmann, D., & Strobel, M. (2004). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *American Economic Review*, 94, 857–869.
- Eriksson, K., Strimling, P., Andersson, P. A., & Lindholm, T. (2017). Costly punishment in the ultimatum game evokes moral concern, in particular when frames as payoff reduction. *Journal of Experimental Social Psychology*, 69, 59–64.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114, 817–868.
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162, 1243–1248.
- Henrich, J. (2006). Cooperation, punishment, and the evolution of human institutions. *Science*, 311(5769), 60–61.
- Hoffman, M., Yoeli, E., & Nowak, M. A. (2015). Cooperate without looking: Why we care what people think and not just what they do. *Proceedings of the National Academy of Sciences*, 112(6), 1727–1732.
- Horton, J. J., Rand, D. G., & Zeckhauser, R. J. (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental Economics*, 14, 399–425.
- Huang, L., Lei, W., Xu, F., Yu, L., & Shi, F. (2019). Choosing an equitable or efficient option: A distribution dilemma. *Social Behavior & Personality*, 47, Article e8559.
- Human, S., & Capraro, V. (2020). *The effect of nudging personal and injunctive norms on the trade-off between objective equality and efficiency* (Available at SSRN 3703684).
- Jordan, J. J., Hoffman, M., Bloom, P., & Rand, D. G. (2016). Third-party punishment as a costly signal of trustworthiness. *Nature*, 530, 473–476.
- Jordan, J. J., Hoffman, M., Nowak, M. A., & Rand, D. G. (2016). Uncalculating cooperation is used to signal trustworthiness. *Proceedings of the National Academy of Sciences*, 113, 8658–8663.
- Jordan, J. J., & Rand, D. G. (2020). Signaling when no one is watching: A reputation heuristics account of outrage and punishment in one-shot anonymous interactions. *Journal of Personality and Social Psychology*, 118(1), 57–88.
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Krupka, E. L., & Weber, R. A. (2013). Identifying social norms using coordination games: Why does dictator game sharing vary? *Journal of the European Economic Association*, 11, 495–524.
- Lacetera, N., & Macis, M. (2010). Social image concerns and prosocial behavior: Field evidence from a nonlinear incentive scheme. *Journal of Economic Behavior & Organization*, 76, 225–237.
- Levine, D. K. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1, 593–622.
- Mazar, N., Amir, O., & Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research*, 45(6), 633–644.
- Milinski, M., Semmann, D., & Krambeck, H.-J. (2002). Reputation helps solve the ‘tragedy of the commons’. *Nature*, 415(6870), 424–426.
- Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, 314, 1560–1563.
- Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, 437(7063), 1291–1298.
- Paolacci, G., & Chandler, J. (2014). Inside the Turk: Understanding mechanical Turk as a participant pool. *Current Directions in Psychological Science*, 23, 184–188.
- Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon mechanical Turk. *Judgment and Decision making*, 5, 411–419.
- Perc, M., Jordan, J. J., Rand, D. G., Wang, Z., Boccaletti, S., & Szolnoki, A. (2017). Statistical physics of human cooperation. *Physics Reports*, 687, 1–51.

- Rand, D., & Nowak, M. (2013). Human cooperation. *Trends in Cognitive Sciences*, 17(8), 413–425.
- Rand, D. G., Tomlin, D., Bear, A., Ludvig, E. A., & Cohen, J. D. (2017). Cyclical population dynamics of automatic versus controlled processing: An evolutionary pendulum. *Psychological Review*, 124, 626–642.
- Rege, M., & Telle, K. (2004). The impact of social approval and framing on cooperation in public good situations. *Journal of Public Economics*, 88(7–8), 1625–1644.
- Tappin, B. M., & Capraro, V. (2018). Doing good vs. avoiding bad in prosocial choice: A refined test and extension of the morality preference hypothesis. *Journal of Experimental Social Psychology*, 79, 64–70.
- Trivers, R. L. (1971). The evolution of reciprocal altruism. *The Quarterly Review of Biology*, 46, 35–57.
- Van Bavel, J. J., Baicker, K., Boggio, P., Capraro, V., Cichocka, A., Cikara, M., ... Willer, R. (2020). Using social and behavioural science to support COVID-19 pandemic response. *Nature Human Behaviour*, 4, 460–471.