Stationary distribution of moral strategies in a population

Jeroen M. van Baar^{1,2*}, Felix Klaassen², Filippo Ricci^{2,3}, Luke J. Chang⁴, Alan G. Sanfey^{2,5}

Author affiliations

¹ Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, Providence, RI 02912, United States

² Donders Institute for Brain, Cognition, and Behavior, Radboud University, 6525 EN, Nijmegen, The Netherlands

³ Utrecht University School of Economics, Utrecht University, 3584 EC, Utrecht, The Netherlands

⁴ Department of Psychological and Brain Sciences, Dartmouth College, Hanover, NH 03755, United States

⁵ Behavioral Science Institute, Radboud University, 6525 HR, Nijmegen, The Netherlands

* Corresponding author at 149 Metcalf Hall, 190 Thayer St, Providence, RI 02912, United States

Email: jvb@brown.edu

Phone: +1-401-863-2998

ORCIDs

0000-0001-8567-7931 (J.v.B.)

0000-0002-6621-8120 (L.C.)

Keywords

Social decision-making, reciprocity, motivation, individual differences.

Author Contributions

J.v.B. and A.S. designed research. J.v.B., F.K., and F.R. performed research. J.v.B., F.K., F.R., and L.C. analyzed data. J.v.B. and A.S. wrote the paper. All authors edited the manuscript.

Abstract

Evolutionary models show that human cooperation can arise through direct reciprocity relationships. However, it remains largely unclear which psychological mechanisms may proximally motivate an individual to reciprocate. Recent evidence demonstrates that psychological motives for reciprocal choices (i.e., moral strategies) differ between individuals, which raises the question whether these differences have a stationary distribution in a population or are rather an artifact of the experimental task. Here, we combine data from three independent studies and participant samples to find that the relative prevalence of different moral strategies is highly stable across these datasets. Furthermore, the distribution of moral strategies is relatively unaffected by changes to the salient features of the experimental paradigm. Finally, the moral strategy classification assigned by our computational modeling analysis corresponds to the participants' own subjective experience of their psychological decision process, and no existing models of social preference can account for the observed individual differences in moral strategies. This research supports the view that social decision-making is not just regulated by individual differences in 'pro-social' versus 'pro-self' tendencies, but also by trait-like differences across several alternative pro-social motives, whose distribution in a population is stationary.

Introduction

Cooperation is central to the evolutionary success of humans and many other organisms¹. One important mechanism for cooperation is reciprocity, defined as situations in which individuals take turns incurring a small personal loss in order to benefit the other². Reciprocity is regulated by psychological processes including guilt, gratitude, and aggression; these processes can be favored by natural selection given the right environmental conditions³. Furthermore, several social preferences have been posited as potential motivations for reciprocal decisions. For example, according to inequity aversion (IA), individuals are motivated to minimize unequal payoffs between themselves and their interaction partner⁴⁻⁷. Alternatively, participants may be motivated by guilt aversion (GA) to avoid feelings of guilt associated with disappointing an interaction partner⁸⁻¹⁰. Recent evidence shows that different people may be motivated by qualitatively different sets of these social preferences (i.e., moral strategies), and some individuals may exhibit context-dependent preferences such as moral opportunism (MO)¹¹. This raises the question whether the distribution of moral strategies across people in a given population is stable and predictable, or rather changes with time and decision task. This is an important question, since the distribution of social preferences in a population can determine whether an intervention to increase prosocial action backfires or succeeds¹²⁻¹⁵.

To address this question, we test two competing hypotheses about the distribution of moral strategies observed in experiments across three studies. The *population property hypothesis* holds that the experimentally observed distribution of moral strategies reflects a stable property of the population sampled. This hypothesis follows from a large body of work on cultural differences in moral judgment, ranging from studies on value differences between tech employees around the world¹⁶ to work on intercultural differences in moral foundations¹⁷⁻¹⁹. If the population property hypothesis is true, we should find a highly similar distribution of moral strategies (inequity-averse participants, guilt-averse participants, and so on) *each time* we run an experiment in a new sample of participants drawn from the same population. In contrast, the *strategy salience hypothesis* posits that the observed distribution of social preferences is primarily a consequence of contextual effects, such as the particular framing of a study. If this hypothesis is true, we should observe significant changes in the distribution of social preferences across instantiations of economic games, each with a different framing and strategy emphasis. Prior economic game experiments do not inform these hypotheses, for although they demonstrate differences in social choice behavior between populations, they remain agnostic about the underlying social preferences, and populations are rarely probed repeatedly to examine the stability of social choice patterns (e.g.²⁰⁻²²).

We test these hypotheses using an experimental task known as the Hidden Multiplier Trust Game (HMTG)¹¹, which is based on the standard Trust Game or Investment Game²³. In a Trust Game, Player 1 (the Investor) can choose to invest any proportion of a \$10 endowment in Player 2 (the Trustee), retaining the remainder. The investment is multiplied x4 before being sent to the Trustee, who can then freely decide how much (if any) of the multiplied amount to return to the Investor. The Trustee commonly returns about 50% of the multiplied investment back to the Investor²⁴, which is in line with both the inequity

aversion and guilt aversion models ^{11,25} and thus introduces a confound in the interpretation of Trustee reciprocity decisions. To solve this problem, the Hidden Multiplier Trust Game multiplies the investment by x2 or x6 instead of x4 on half the trials. The Investor however is led to believe that their investment is always multiplied by x4, while the Trustee knows the true multiplier and knows about the Investor's ignorance. In recent work, this paradigm has revealed clear individual differences across participants playing as Trustees in the degree to which inequity- and guilt-averse preferences influenced their reciprocity decisions ^{11,26}. Computational modeling showed that about 40% of Trustee participants are mainly motivated by inequity aversion (always returning about half the multiplied investment), 10% are mainly guilt-averse (always returning what the Investor expects, insensitive to the changing multiplier), 40% are morally opportunistic (switching between inequity-averse and guilt-averse choices depending on the task context, always choosing the most financially advantageous of the two), and 10% are greedy (they return little to no money to the Investor) ¹¹.

Two properties of the Hidden Multiplier Trust Game allow us to probe whether this distribution of social preference types is a stable feature of the population or an artifact of the experimental task. First, the *task incentives* of the HMTG place a large financial pressure on guilt-averse Trustees because these players need to sacrifice all their game tokens in the x2 condition to meet the Investor's expectations. In study 1, therefore, we limit the financial pressure on guilt-averse participants by administering the HMTG twice to the same participants, once with the original x2/x4/x6 multipliers and once using multipliers x4, x6, and x8 (investor believes x6), testing how stable participants' strategies are under changing incentives. Second, the *salient strategy* in the HMTG is inequity aversion, as the experiment manipulates the true multiplier by which each investment is multiplied before being sent to the Trustee, which yields a change in inequity-averse behavior while guilt-averse choices stay the same. In study 2, therefore, we presented new participants with a 'flipped' version of the task, the False-Belief Multiplier Trust Game (FBMTG), where the Investor believes the multiplier is x2, x4, or x6, but the Trustee learns that the true multiplier is always x4. Here guilt aversion is the salient strategy as the task manipulates the expectations of the Investor, and we test whether this raises the prevalence of guilt aversion and lowers the prevalence of inequity aversion as predicted by the strategy salience hypothesis.

Across both studies, we employ an integrative model of social preferences to quantitatively map the moral strategies of our participants. In this Moral Strategy model (eq. 1; ¹¹), two free parameters (Θ ; and Φ) determine the influence of three sources of utility—payoff (π), guilt, and inequity—on decisions. For a detailed description of the model, see Methods.

$$U_2(S_2) = \Theta * \pi_2 - (1 - \Theta) * \min(Guilt_2 + \Phi, Inequity_2 - \Phi) \quad (1)$$

Results

Computational modeling

Different participants showed markedly different patterns of reciprocity behavior in the Hidden Multiplier Trust Game (illustrated in fig. 1 for study 1). To formally test whether these differences reflected between-subject differences in moral decision strategy rather than noise around a single strategy shared across all participants, we compared the model fit of the three unitary social preference models (greed, guilt aversion, and inequity aversion) to that of our Moral Strategy model ¹¹. This latter approach integrates these social preferences into one model and allows the individual weights to vary across participants.

In study 1, we fit all four models to the participants' behavior separately in each of the two HMTG blocks (2/x4/x6 and x4/x6/x8). In the x2/x4/x6 block, model fit was significantly better for the Moral Strategy (MS) model than for the greed model, the guilt aversion model, and the inequity aversion model (paired-samples t-test on AIC between MS model and best other model, i.e. MS versus IA: $t(93) = -2.65$, $P = 0.0096$; fig. 2A). Although the improved model fit for the MS model relative to IA was not significant in the x4/x6/x8 context (MS versus IA: $t(93) = -1.31$, $P = 0.194$), mean model performance across the two contexts was significantly better for MS (paired-samples t-test on mean subject AICs across contexts, IA versus MS: $t(93) = 2.07$, $P = 0.042$).

In study 2, model fit was significantly better for the Moral Strategy model compared to the unitary preference models (i.e., greed, guilt aversion, and inequity aversion models) (fig. 2B) (paired-samples t-test on AIC between IA and MS models: $t(52) = 3.31$, $P = 0.002$). We thus replicate our earlier results ¹¹, showing that a set of qualitatively different moral strategies can better capture the behavior in the HMTG than any unitary social preference model by itself, even when penalizing for increased model complexity. Posterior predictive checks on four representative participants (fig. 1) confirm that the Moral Strategy model better captured the key differences between participants than did either of the guilt aversion and inequity aversion models.

Fig 1. Left panels: Task behavior for four example participants in study 1, block x2/x4/x6 (A) and x4/x6/x8 (B). Other panels: Model predictions for the GA, IA, and MS (moral strategy) models, fitted to these behavioral data, are shown also. This posterior predictive check reveals that the MS model was best at capturing the four predicted behavioral strategies that are exemplified by these four example participants. The superior fit of the MS model is confirmed quantitatively in formal model comparison (see fig. 2A).

Fig. 2. Model comparisons for the moral strategy (MS), inequity aversion (IA), guilt aversion (GA) and greed (GR) models. A) Model comparisons across the two contexts of study 1. B) Model comparisons in study 2. AIC: Akaike Information Criterion. **: $P < 0.01$ (uncorrected). Error bars represent bootstrapped 95% confidence intervals.

Determining participants' moral strategies

One notable advantage of the Moral Strategy model is that by fitting this model to each participant's behavior, we obtained a pair of parameters (theta & phi) that best captured that particular participant's reciprocity motives. We can then classify each participant's moral strategy based on their position in the theta-phi parameter space of the model. For each of the four theoretically predicted moral strategies (GR, GA, IA, MO), a zone can be identified where this moral strategy is best represented, by use of the model-driven clustering procedure described in ¹¹. In brief, this clustering method consists of first simulating task behavior across equidistant theta-phi coordinates in the parameter space, and then using a hierarchical clustering algorithm to group these simulations based on pairwise similarity in simulated task behavior (for details see Methods). Of note, this method does not rely on any participant data, and can therefore be flexibly applied across experiments with different task parameters (e.g. different sets of multipliers).

Figure 3 shows the position of each participant in the theta-phi model parameter space across both studies and sets of multipliers, overlaid on the colored zones that represent the four moral strategies. Participants are labeled by moral strategy based on their position in one of the four strategy zones, and

participant behavior for each cluster is plotted in Supplemental Figure 1 as a qualitative confirmation of clustering accuracy.

Figure 3. Model parameter space in study 1 (A) and 2 (B). The space contains the four hypothesized moral strategies at different zones of parameter combinations (colored zones), with best-fitting parameters added in for each participant (scattered points). Circles: greed; stars: guilt aversion; squares: inequity aversion; diamonds: moral opportunism.

Can financial pressure drive guilt-averse participants to moral opportunism?

To examine the stability of participants' moral strategies across blocks in study 1, we quantified the test-retest reliability of model parameters theta and phi. It should be noted that small changes in these parameters are expected across blocks, since the boundary between strategies in the parameter space changes slightly as a function of multiplier set (x2/x4/x6 versus x4/x6/x8, see figure 3AB). Therefore, we also computed test-retest reliability of the distance between each pair of participants in the theta-phi parameter space. If participants employ stable motives across blocks, these inter-participant distances should remain stable even if small shifts in theta/phi occur for each participant. For test-retest reliability, we computed the Pearson correlation for theta, phi, and inter-participant distance between blocks x2/x4/x6 and x4/x6/x8. All three metrics yielded very high reliability across contexts (theta: $r = 0.95$, $p < 0.001$; phi: $r = 0.66$, $p < 0.001$; inter-subject distance: $r = 0.86$, $p < 0.001$; fig. 4A-C), which suggests that participant behavior was relatively insensitive to the multiplier manipulation, which supports the population property hypothesis but not the strategy salience hypothesis.

As a more precise test of the strategy salience hypothesis, we hypothesized that the financial pressure of the x2 condition of the HMTG could push otherwise guilt-averse players towards a strategy of moral opportunism. If this interpretation were true, we would expect participants who were morally opportunistic in the x2/x4/x6 block to exhibit a guilt-averse behavior pattern in x4/x6/x8. We tested this prediction using a Stuart-Maxwell test for paired (repeated-measures) frequency sets, which revealed no significantly different distribution of strategies between the two multiplier contexts ($\chi^2(3) = 1.48$, $P = 0.69$;

fig 4D). A more specific Stuart-Maxwell test for switching between MO and GA revealed no significant change in relative frequency of these two strategies between the two contexts ($\chi^2(1) = 0.33$, $p = 0.56$). On the contrary, as fig. 4D shows, most participants were consistent in their moral strategy between blocks x2/x4/x6 and x4/x6/x8 (the highest values lie on the diagonal), and only 2.1% of moral opportunists in x2/x4/x6 were guilt-averse in x4/x6/x8. This suggests that the financial pressure imposed by the x2 condition did not specifically shape the distribution of moral strategies in the x2/x4/x6 context, which supports the population property hypothesis over the strategy salience hypothesis.

Figure 4. Moral strategy is stable across the x2/x4/x6 and x4/x6/x8 context of study 1. A-B) Test-retest reliability of model parameters theta and phi between the two contexts is high. C) Test-retest reliability of the inter-subject distance in the model parameter space is high. Inter-subject distance captures the geometry of similarity between participants in theta-phi space, which is useful since small changes in parameter values are expected between the x2/x4/x6 and x4/x6/x8 blocks due to the changing multiplier magnitude. D) Categorical shifts in moral strategy between the two contexts are rare (highest values on the diagonal).

Does task salience affect moral strategy?

Study 2 was designed to make guilt aversion more salient than in study 1 by manipulating the expectations of the Investors, and should conversely make inequity aversion less salient than in study 1. To test these predictions of the strategy salience hypothesis we compared the relative prevalence of the four moral strategies between study 2 and the x2/x4/x6 block of study 1 using a chi-square test for independent samples. The results are plotted in figure 5, alongside the distribution of moral strategies found in our prior work ¹¹. Overall, the relative prevalence of the four strategies was significantly different between study 1 and study 2 ($X^2(3) = 8.43$, $P = 0.038$). At first glance, this effect appears to be driven by an increased prevalence of GA in study 2 as compared to study 1 (fig. 5): in study 2 GA is no longer the least popular moral strategy, but the second least popular. However, this increase in GA prevalence was not accompanied by a decrease in IA prevalence, as IA was about equally common in study 1 as in study 2. Indeed, when testing the between-study difference in relative prevalence of each pairwise combination of two strategies, the only significant change was in the relative prevalence of GA and GR ($X^2(1) = 6.40$, P (uncorrected) = 0.011), which may be a false positive since this effect does not hold when Bonferroni-correcting for multiple pairwise tests (P (corrected) = 0.069). IA and MO were just as dominant in study 2 as in study 1. Given this mixed evidence, we additionally compared the strategy prevalence in study 2 to that of our prior work on the HMTG ¹¹. Between these studies, no significant change in strategy prevalence was found ($X^2(3) = 6.54$, $P = 0.088$), and none of the pairwise prevalence shifts were significant (all P (uncorrected) > 0.05). Taken together, these results provide some support for the strategy salience hypothesis, but also highlight the striking consistency of moral strategies across three different datasets, collected at different times, in different experimental settings (fMRI, behavioral), and, crucially, with entirely different salient features in the experimental paradigms.

Figure 5. The relative prevalence of the four moral strategies in prior work (fMRI dataset; Van Baar et al., 2019), study 1 (x2-x4-x6 context), and study 2. HMTG: Hidden Multiplier Trust Game. FBMTG: False-Belief Multiplier Trust Game. * Difference in moral strategy distribution is significant at $P < 0.05$; n.s. not significant.

Construct validity of moral strategy

The combined results of study 1 and 2 suggest that the observed distribution of moral strategies is a stable property of the population from which we sampled. This interpretation naturally raises the question whether the moral strategy labels we apply to participants in fact correspond to the participant's own experience of their psychological decision process (i.e., construct validity). We tested this by asking all participants in the two studies to self-report on how they made their decisions in the HMTG and FBMTG, suggesting four potential strategies and leaving room for other strategies in an open-ended answer box. To maximize our statistical power for this follow-up analysis we pooled data across studies 1 and 2. To avoid including participants who were unstable in their moral strategy, we kept only the 74 participants from study 1 who exhibited the same moral strategy category across the two multiplier blocks. The total pooled n was therefore $74 + 55 = 129$.

As predicted, there were significant differences in the weight assigned by the four groups (GR, GA, IA, MO) to the first three of the suggested strategies ('Keep': $F(3,125) = 52.51$, $p < 0.001$; '50-50': $F(3,125) =$

33.49, $P < 0.001$; 'Expectation': $F(3,125) = 13.63$, $p < 0.001$) (fig. 6). To better understand these findings, we tested whether each of these three suggested strategies was self-reported as most important by the associated moral strategy group. This was indeed the case: Inequity-averse participants assigned significantly higher weight to '50-50' than the other three groups (IA-GA: $P = 0.006$; all other IA pairs $P < 0.001$); greedy participants assigned significantly higher weight to 'Keep' than the other three groups (all $P < 0.001$); and guilt-averse participants rated the 'Expectation' strategy as more important than all other groups, although this effect was not significant for GA-MO ($P = 0.16$; all other GA pairs $P < 0.01$). Taken together, these results lend construct validity to our model-based strategy interpretations, suggesting that we have appropriately labeled the reciprocity motives present in our participants.

Fig. 6. Self-reported importance of the three most prominent decision strategies for participants in the four model-derived moral strategy groups. Points are mean over group; solid lines represent standard error of the mean; dotted lines are perpendicular projections from the mean points onto the axes.

Determinants of moral strategy

Can any other commonly used measures of social decision strategy account for the individual differences reported in this paper? To answer this question, we compared the model-derived moral strategies to participants' scores on three other measures: Social Value Orientation (SVO; ²⁷), dispositional greed ²⁸, and trait guilt ²⁹. If model parameter theta (the 'greed parameter') accurately captures the greed aspect of moral strategy, we would predict that, across participants, this parameter is

negatively correlated with the social value orientation (SVO) score, which is known to predict prosociality in a range of laboratory tasks^{27,30}. In line with this expectation, we found a strong and significant negative correlation between theta and SVO ($r = -0.68$, $p < 0.001$), such that higher SVO scores were related to lower theta values. We also found a significant, though weaker, positive correlation between theta and dispositional greed (DG) ($r = 0.21$, $p = 0.017$). Theta was not correlated to trait guilt, as expected. Due to a significant negative correlation between phi and theta ($r = -0.38$, $P < 0.001$), raw phi values also correlated with SVO, but in the opposite direction from the theta relationships (phi with SVO: $r = 0.28$, $P = 0.0011$). To control for this confound, we residualized phi with respect to theta (i.e. took the residual of the regression of phi onto theta); residualized phi was not significantly correlated with either SVO ($P = 0.74$) or DG ($P = 0.22$). Surprisingly, residualized phi was also not associated with trait guilt ($P = 0.61$). These results suggest that previously proposed measures of social preference (social value orientation and greed) capture the same trait as our model's theta parameter, but that parameter phi captures previously unexplained variance in social choice behavior. Therefore, Hidden Multiplier games offer a novel approach to both examine actual behavior in an interactive setting, as well as offering a unique assay for directly measuring the motives underlying reciprocity behavior.

Discussion

In this work, we demonstrate that (a) different people follow qualitatively different moral strategies when making reciprocity decisions in a modified Trust Game, (b) these strategies are stable over time within people, (c) the distribution of strategies across datasets is stable, and (d) the salience of a specific moral strategy only weakly affected the prevalence of that strategy in the data. We further illustrate that the moral strategies, derived from computational utility models grounded in economic theory, align with participants' own experience of their psychological decision process, which lends construct validity to our observations. Taken together, these findings strongly support the population property hypothesis over the strategy salience hypothesis, suggesting that social preferences are distributed in a population in a stationary manner. In other words, our results show that social decision-making is not just regulated by individual differences in 'pro-social' versus 'pro-self' tendencies, but also by trait-like differences across several alternative pro-social motives, whose distribution in a population is stationary.

Our observations raise two questions for further research on the prevalence of moral strategies in a population. First, it remains unclear precisely which mechanisms buffer against changes in the moral strategy distribution. One possibility is that participants randomly choose a moral strategy to follow when they enter into a social decision-making experiment, and that the probability with which they choose a given strategy (IA, MO, etc.) is somehow fixed in a population, for example via culturally distributed associations between tasks and social norms. An alternative possibility is that a person's moral strategy in laboratory tasks reflects a stable 'moral phenotype' that drives their social choice behavior across contexts. This interpretation is in line with other evidence suggesting that social decision-making strategies are linked to stable traits including interpersonal attachment style³¹ and political preference^{19,32}. Yet, this interpretation goes beyond prior work in suggesting that people are stable not only in their

behavioral tendencies for pro-social or selfish choices, but also in the psychological processes by which they make their moral decisions. Future work could usefully examine the development of moral phenotypes at the motivational (rather than behavioral) level, for example by exposing the same group of participants to different social decision tasks separated by long time intervals (in the order of months) and measuring cross-task predictive power. Such work should take existing data on the mismatch between laboratory task decisions and real-world moral behavior into account ³³. Follow-up work could also consider testing explicitly for other motives such as reciprocity motives (Falk et al., 2003).

Secondly, it is an exciting open question which features of the population might give rise to a particular distribution of moral strategies. For example, population-wide trends in personality traits or in cultural values might go hand-in-hand with moral strategies, and we hope future work will illuminate these possibilities. We believe that such work should incorporate a computational modeling approach to moral strategies in social decision-making, for several reasons. First, our Moral Strategy model explained behavioral variance unaccounted for by prior modeling approaches: although widely used metrics of greed and prosociality strongly predicted the greed parameter of our computational model (which lends measurement validity to the model), the variance explained by the social preference dimension of our model (parameter ϕ) was not captured by any other measure that we collected here. Second, the participants in the different moral strategy groups—which were derived from clustering of the computational model’s parameter space—self-reported different motives for Trust Game reciprocity, which lends construct validity to our model. Third, we show that a computational modeling approach can allow for precise quantification of motives in social decision-making, which can greatly help support the aforementioned goal of discovering personality and cultural correlates of moral strategies.

We observed a weak effect of task salience on moral strategy. This is somewhat surprising, as demand effects in social psychology are common ^{34,35} and suggest that participants should track the most clearly manipulated component of a task in their behavior. Nevertheless, the increase in prevalence of GA relative to GR between studies 1 and 2 may indicate that demand effects may play a modest role in some motives for reciprocity, while other reciprocity motives (e.g. inequity aversion) are always salient, regardless of task manipulation.

Interestingly, the individual difference measure of ‘trait guilt’ in these data (part of the Guilt Inventory as described by ²⁹) did not predict guilt-averse preferences in the HMTG. This lack of relationship can be understood by investigating what is meant by ‘guilt’ in these two different contexts. The Guilt Inventory questionnaire items loading on ‘trait guilt’ relate to the guilt one feels individually when thinking back to a past misdeed (e.g. one item reads “I have made a lot of mistakes in my life”), as well as guilt related to disappointing authority (e.g. “My parents were very strict with me”). Guilt aversion as defined in our study and by ⁸, instead describes an interpersonal emotion ³⁶ that results from disappointing an equal interaction partner. These two types of guilt thus point to distinct psychological experiences. Importantly, ‘guilt’ as operationalized in the Guilt Inventory could follow from behaving unfairly just as well as from disappointing an interaction partner, while our definition of guilt aversion is specific to the latter. Future work could more clearly tease apart these two components, perhaps by developing a questionnaire measure for

interpersonal guilt, and could also clarify the relationship between these types of guilt on the one hand, and regret and shame on the other.

This research has potential implications for the design of government interventions intended to foster cooperative behavior of citizens in a society, as the effectiveness of such interventions depends on people's endogenous social preferences¹³. For example, the interventions required to reduce tax evasion in *guilt-averse individuals* may be very different than those required for *moral opportunists*—specifically, to motivate moral opportunists one likely needs a multi-pronged approach where the opportunists' entire portfolio of moral reasoning strategies is simultaneously influenced. Designers of interventions should therefore heed the moral motives underlying our social decisions, with the good news being that the distribution of moral strategies in a population appears to be stable over time and tasks. Similarly, financial incentives can erode intrinsic prosocial motives for cooperation through 'motivational crowding-out'^{15,37}. For instance, an incentive can backfire by changing the framing of a moral decision from ethical to self-interested (e.g.³⁸). Since such framing effects depend on how people frame a cooperative decision to begin with, our observation that the subjective experience of reciprocity decisions differs between moral strategy types is significant. Motivational crowding-out may also occur when an incentive signals that *other people* may be mainly motivated by self-interest^{15,37}, for instance when the incentive is so large that cooperation is no longer a costly signal of one's prosocial intentions. The effectiveness of an incentive thus depends on the beliefs people have about other decision-makers' moral motives, and future research should address whether these beliefs accord with the true distribution of moral strategies.

Our findings caution against interpreting social choice behavior through the lens of a single moral motive, and suggest that individual differences in moral choice are a crucial part of understanding decisions made in these contexts. The traditional Trust Game could not have uncovered the moral strategy variance observed here: in the x4 condition of the HMTG the behavior predicted by guilt aversion, inequity aversion, and moral opportunism is the same. In general, therefore, we argue for more fine-grained analysis of the moral motives underlying the social choice behavior. By formally modeling different motives and teasing them apart using innovative experimental tasks, we can begin to better understand idiosyncratic laboratory behavior as well as the social decisions that define our societies.

Materials and methods

Study 1

The Hidden Multiplier Trust Game

Both studies in the present paper are developed from prior work using the novel Hidden Multiplier Trust Game (HMTG;¹¹). In each trial of this game, an anonymous Investor gets the chance to invest any number of 10 game tokens with an anonymous Trustee, retaining the remainder. These are all 'single-shot' games, that is, the Investor and Trustee interact only once. As in a standard Trust Game²³, the invested amount is increased by a multiplier before being transferred to the Trustee. The Trustee then

gets the chance to return any number of tokens from the multiplied amount to the Investor, but, importantly, does not have to do so. The tokens are redeemed for money at the end of the experiment (exchange rate: €0.40 per token).

In the standard Trust Game, the multiplier is x4, and both the Investor and the Trustee have full knowledge of this. In the Hidden Multiplier Trust Game, however, the multiplier is sometimes x2 (25% of rounds) or x6 (25%) instead. Crucially, the Investor believes the multiplier is always x4, whereas the Trustee knows both about the multiplier used, as well as about the Investor's ignorance as to the actual multiplier. In the HMTG, therefore, the Trustee sometimes has more or fewer tokens than the Investor believes, and the expectations of the Investor are uncoupled from the true amount of money received by the Trustee.

Whereas guilt aversion and inequity aversion models make the same behavioral predictions in the classic Trust Game, these predictions diverge in the HMTG. Specifically, a guilt-averse Trustee, motivated to satisfy the expectations of the Investor ⁸, should send the same amount of money back to the Investor in each of the three multiplier conditions, as the Investor's expectations do not change across these conditions. An inequity-averse Trustee, on the other hand, is motivated to provide an even split of all available tokens in the game ^{4,5}, and should therefore return more money in the x6 condition than in the x4 condition, and more in x4 than in x2, that is, they should be motivated by the actual amount of tokens they possess. A morally opportunistic Trustee combines these two decision strategies adaptively, returning enough to satisfy the Investor's expectations in the x6 condition, while creating an even split in the x2 condition. A greedy Trustee, finally, should not send back any tokens. These four 'moral strategies' are formalized in the computational model below.

Participants

104 adult participants were recruited from the Nijmegen student population using an online recruitment tool. People were excluded from participation if they had taken part in an experiment on the Hidden Multiplier Trust Game before, if they were students of economics, if they were not native speakers of Dutch, if they were pregnant, and/or if they suffered from claustrophobia.

Task

All participants played two blocks of the Hidden Multiplier Trust Game ¹¹, each with 80 trials. One of the two blocks of 80 rounds of the HMTG was played using the multiplier set of [x2, x4, x6], where the Investor believed the multiplier was always x4. The other block was played with multiplier set [x4, x6, x8], and here the participants were instructed that the Investors always believed the multiplier was x6. The order of the two blocks was counterbalanced across participants.

Procedure

The experiment consisted of a single session. Upon arrival at the laboratory, each participant was seated in a closed-off behavioral lab space, where they were screened for participation, were informed about the

study, and gave written informed consent. Next, they received written general instructions about the standard Trust Game. To avoid biasing game behavior, the Trust Game was always referred to as 'Investment game', the Investor as 'player A', and the Trustee as 'player B'. The multiplier in these standard instructions was the *Investor-believed multiplier* of the block in which the participant started (i.e. x4 for x2/x4/x6 block first, and x6 for x4/x6/x8 block first). Importantly, these instructions contained no information about the 'hidden multiplier' aspect of the task, and were meant to inform the participants about the general structure of this game. In these instructions, the participants were asked how many of 10 game tokens they would themselves invest in an anonymous Trustee, if they were the Investor (and the multiplier was the base multiplier of x4 or x6). They were asked for consent to use this investment in a future experiment on the Trust Game, and were informed that if their investment was used in the future, they would receive additional payment based on the amount sent back to them by the future Trustee ('player B'). We additionally asked for permission to take their portrait photo, which was to be used alongside their investment in the potential future experiment. All these questions were added to the instructions to ensure that the investments made here by the participants were sincere, and to ensure that the participants believed that they were interacting with real investors (i.e. people who participated at an earlier time). In reality, the participants played with pre-recorded investment data from a single-shot Trust Game with multiplier x4¹¹. This instruction step additionally allowed us to generate a new set of true investments for future experimental use.

After reading the general Trust Game instructions, the participants read additional instructions pertaining to the 'hidden multiplier' element of the Hidden Multiplier Trust Game. Here, we informed them that the Investors they would encounter always believed that the multiplier was fixed (at x4 for participants who started with the x2/x4/x6 multiplier block, and at x6 for x4/x6/x8), but that the actual multiplier would sometimes be lower (x2 or x4) or higher (x6 or x8). After testing the participants on their understanding of these instructions in a short quiz, they played 80 single-shot, anonymous rounds of the Hidden Multiplier Trust Game, always in the role of Trustee, in the context they were assigned to first (x2/x4/x6 or x4/x6/x8).

On each round, the Investor's participant number was presented alongside a blurred photo of a face, which was added to strengthen the belief that the Trustees were playing with real Investors. The photos were taken from the Radboud Faces Database³⁹ and blurred heavily in Matlab to ensure that participants could not recognize the gender or identity of the faces. On the next screen, the participants saw how many tokens the Investor had invested in them, and by which factor this investment was multiplied. The participants were also reminded that the Investor always believed the multiplier was x4. On the final screen, the participants could indicate how many tokens from the multiplied amount they wanted to return to the Investor; crucially, they could also choose to keep it all. There was a short break after 40 trials.

After playing 80 rounds in the first context, the participants were allowed to take a self-paced break while remaining in the behavioral lab space, and received instructions for the second context. These short instructions highlighted the fact that the multiplier set would change, and that the new Investors for this second context always believed the multiplier was x6 (in x4/x6/x8) or x4 (in x2/x4/x6). Thus, all

participants played 80 rounds of the HMTG twice, once in x2/x4/x6 and x4/x6/x8, with context order counterbalanced between participants, and they were informed that they played with a unique anonymous Investor on each of the 160 rounds of this experiment. The investment sets were identical between the two contexts. The participants were informed that one trial out of the 160 would be randomly selected to be paid out both to them and to the corresponding Investor.

Since the multiplier in the HMTG determines with how many tokens the Trustee gets to play, it also determines the behavioral predictions of the inequity aversion model: under the raised multiplier condition (x6 or x8) the amount sent back to the Investor by the Trustee ought to be higher than under the base multiplier (x4/x6) or the lowered multiplier (x2/x4). Since the Investor always believes the multiplier to take the base value (x4 in x2/x4/x6 and x6 in x4/x6/x8), however, the predictions of the guilt aversion model are identical across multiplier conditions. This allows us to deduce from each Trustee's behavior which moral strategy they followed in their reciprocity decisions (see *Computational modeling*).

Additional measures

After finishing the HMTG, participants completed a computerized version of the social value orientation (SVO; ²⁷) task, which was based on the 'slider measure' version of the SVO ³⁰. Then, they completed the dispositional greed scale ²⁸, which measures greediness as a personality trait, and the guilt inventory ²⁹, which yields an individual difference measure of general sensitivity to guilt (trait guilt). Finally, they answered several questions about their own decision-making strategy in the HMTG, dividing 100 points over various descriptions of strategies:

- Inequity aversion ("I wanted to divide the total number of tokens evenly over the Investor and myself")
- Guilt aversion ("I wanted to give the number of tokens that the Investor expected to receive")
- Greed ("I wanted to keep as much of the investment as possible for myself")
- Altruism ("I wanted to return as much as possible of the investment")
- Other ("Other, describe and indicate number of points")

Finally, each participant was asked whether they believed an Investor would be happy to be paid based on the participant's reciprocity decisions. If at this point, or at any other time, the participant indicated that they did not believe the Investors were real other participants, their data were excluded from the analysis. Based on these criteria, 2 participants were excluded from analysis, leaving a total of 102 participants (50 with the 2-4-6 context first, 52 with 4-6-8 first) (32 men; mean age 22.3 ± 2.7 y).

Computational modeling

We used the Moral Strategy model ¹¹ to infer the Trustees' motives based on their reciprocity behavior in the HMTG. This model (eq. 1) describes the utility derived by the Trustee from their strategy S_2 (i.e. the number of tokens sent back to the Investor) based on monetary gain, guilt, and inequity. Two free parameters (θ , Θ ; and ϕ , Φ) determine the influence of these three sources of utility (see eq. 1 in

main text). In this model, the Trustee's monetary gain is defined as $\pi_2 = (I * M_2 - S_2) / (I * M_2)$, where I represents the Investor's investment and M_2 represents the multiplier known only to the Trustee. Inequity is defined $Inequity_2 = ((I * M_2 - S_2) / (10 - I + I * M_2 - 1/2))^{2.5,11}$, and guilt is defined $Guilt_2 = ((E_2(E_1(S_2)) - S_2) / (E_1(M_1) * I))^{2.8,11,40}$, where $E_2(E_1(S_2))$ represents the Trustee's belief about the Investor's expectations of the Trustee's strategy and $E_1(M_1)$ represents the Investor's belief about the multiplier. To maximize generalizability of the model, we set the Trustee's belief about the Investor's expectations to half of the amount the Investor believes the Trustee has: $E_2(E_1(S_2)) = 1/2 * E_1(M_1) * I$ (for validation, see ¹¹). At each reciprocity decision, either guilt or inequity constitutes the social preference term of the Trustee's utility function. By default, at $\Phi = 0$, the model selects the weakest of the two. This structure explicitly accommodates the 'moral opportunism' behavioral pattern, since it allows the Trustee to disregard guilt (i.e. disappointing the Investor) in the lowered-multiplier context (e.g. x2 in the x2/x4/x6 condition), and to disregard inequity in the raised-multiplier context. As ϕ moves away from 0, however, the model's behavioral output becomes biased towards the guilt aversion ($\phi < 0$) or inequity aversion ($\phi > 0$) prediction. Higher levels of θ (Θ) bias behavior towards greed. At different combinations of θ and ϕ , therefore, the model can accommodate guilt aversion, inequity aversion, moral opportunism, and greed.

In this experiment, we compared the performance of the moral strategy model to that of simpler and previously proposed alternative models: greed (eq. 2), inequity aversion (eq. 3), and guilt aversion (eq. 4). In these models, monetary payoff, guilt, and inequity are defined as in the moral strategy model.

$$U_2(S_2) = \pi_2 \quad (2)$$

$$U_2(S_2) = \pi_2 - \Theta * Inequity_2 \quad (3)$$

$$U_2(S_2) = \pi_2 - \Theta * Guilt_2 \quad (4)$$

Across all models, the predicted strategy for the Trustee was the strategy which yielded maximal utility:

$$\hat{S}_2 = \arg \max_{S_2} U_2(S_2) \quad (5)$$

Model performance was measured and compared using the Akaike Information Criterion (AIC; eq. 6) ^{41,42}, which rewards model fit and penalizes model complexity (number of free parameters). We chose to use AIC over the alternative Bayesian Information Criterion, since AIC is superior to BIC if the true data-generating model is not in the model set ⁴³, which is likely true for the current experiment. Assuming that the model errors are normally distributed, AIC is defined as

$$AIC = n * \ln(SSE/n) + k * 2 \quad (6)$$

where SSE represents the residual sum of squares (i.e. the sum over squared differences between model prediction and actual behavior), n represents the number of observations (trials), and k represents the number of free parameters in the model (theta and/or phi). Participants whose behavior was near-perfectly explained by any model ($SSE < 10$) were excluded from model comparisons, since the logarithm of 0 is undefined.

To avoid obtaining a local rather than global optimum, the computational models were fit to each participant's behavior 1000 times, each time with a different random parameter set established as the starting point for the optimizer. For each participant, the best-fitting parameter set from of all solutions was selected.

We used the moral strategy model results to classify participants as greedy, guilt-averse, inequity-averse, or morally opportunistic respectively, using a similar procedure as in earlier work ¹¹. To this end, we first ran simulations of the moral strategy model for all combinations of investment and multiplier at 10,201 (101 by 101) equidistant points in the model's theta-phi parameter space (theta on the domain [0, 0.5] and phi on [-0.1, 0.1]), and applied hierarchical clustering to these simulations based on their pairwise Euclidean distance in reciprocity behavior. In the simulations, we used as second-order expectations half of the amount the Investor believes the Trustee possesses. For the x2/x4/x6 block, we set the clustering distance cutoff such that five clusters were obtained, and then merged the two clusters whose simulations corresponded to the theoretical moral opportunism pattern. For the x4/x6/x8 context, we set the distance cutoff such that four clusters were immediately obtained. This method ensured that high phi corresponded to inequity aversion, low phi to guilt aversion, phi around zero to moral opportunism, and high theta to greed, all of which are line with the design of the moral strategy model, while letting the clustering algorithm draw the precise boundaries between the corresponding zones of theta-phi coordinates in the model's parameter space. The clustering solution was thus not based on the participant behavior that occurred in our sample, but instead on the theoretical range of behavior that the model could capture. We then located each participant in the parameter space by their best-fitting theta-phi parameter pair, and thus placed them into one of the four moral strategy zones. This procedure yielded, for each participant and each block, a moral strategy label: inequity-averse, guilt-averse, morally opportunistic, and greedy. Importantly, this clustering method precisely mirrored the method applied in our neuroimaging study on the Hidden Multiplier Trust Game ¹¹, as well as the method applied in study 2 here, which means the relative prevalence of the four moral strategies can be directly compared across studies.

Software

Stimuli were presented on a Windows pc running PsychToolBox 3.0.11 (www.psychtoolbox.org) for Matlab 2016a (Mathworks, Natick, MA, USA). Questionnaire data, screening, and debriefing were collected using Castor Electronic Data Capture (www.castoredc.com). Computational model analysis was carried out in Python 2.7, where models were fit to data using the *optimize.least_squares* routine in the *Scipy* package version 1.0.0 ⁴⁴.

Ethical approval

The experimental procedures of both studies were approved by the local ethics review board (CMO Arnhem-Nijmegen, The Netherlands).

Study 2

Participants

57 participants were recruited from the Nijmegen student body using an online recruitment tool. Exclusion criteria were identical to those of study 1 and 2 participants were excluded due to disbelief in the task, leaving 55 participants (14 men; mean age 21.8 ± 2.8 y).

Main task

In this study, participants played an 'inverted' version of the Hidden Multiplier Trust Game, the False-Belief Multiplier Trust Game (FBMTG), always in the role of Trustee. In this variant, the multiplier by which the Investment is increased is always x4, but now the *Investors* have varying beliefs about the multiplier. 50% of the 80 Investors encountered by the Trustee (in single-shot interactions) believed their investment would be multiplied by x4, 25% believe in x2, and 25% in x6. The Trustee knows this, and also knows the multiplier is actually always x4. As in study 1, in 50% of trials, the number of tokens received by the Trustee differs from the number of tokens the Investor believes the Trustee has, which means the reciprocity behavior predicted by the guilt aversion and inequity aversion models again diverges. However, whereas in study 1 (and the original HMTG design) the number of tokens sent to the Trustee changes across conditions, in study 2 the number of tokens the Investor believes the Trustee to have changes across conditions. As in study 1, the FBMTG design also allows for morally opportunistic behavior: a morally opportunistic Trustee could return the Investor's expectation in x2 (GA), keeping the surplus; and could create an even split in x6 (IA), thereby disappointing the Investor.

Procedure

After arriving at the laboratory, the participants were seated in a closed-off behavioral lab space, where they were screened for participation and gave written, informed, consent. They then were presented with the general task instructions. The participants were informed that anonymous other participants ('Player A') had made 'Investment Game' investments in them under varying multipliers (x2, x4 and x6). They learned that they would respond to these investments in the role of Trustee ('Player B') and would be free to choose the number of tokens they wished to send back to each anonymous investor. They were also informed that one of the 80 interactions would be randomly selected to be actually paid out to them and the corresponding Investor. In reality, the investment decisions were pre-programmed identically to study 1, and the selected trial was only financially consequential for the Trustee participant.

After reading these instructions, the participants reported how many of 10 game tokens they would invest if they were the Investor, under each of the three multipliers x2, x4 and x6. For each of these

investments, they also reported how many tokens they would expect back from an anonymous Trustee. Next, as in study 1, we asked permission to take the participant's portrait photo and use this alongside their investment(s) in a potential future study. Finally, the participants were given a page with extra instructions for this experiment, which stated that although the Investors believed in varying multipliers (x2, x4 and x6), the multiplier would in fact always be x4. The participants were quizzed on these instructions, and once they understood started playing 80 rounds of the FBMTG.

On each game round, the participants were presented with the Investor's participant number and a blurred photo of their face, as in study 1. They next read how much the Investor had invested in them, and by how much the Investor believed their investment would be multiplied. They were also reminded that the investment would in fact be multiplied by 4. On the final task screen, the participants indicated how many of the multiplied game tokens they wished to return to the Investor, with the option of keeping all. Halfway through the task, i.e. after 40 trials, the participants could take a (self-paced) break.

Additional measures

After the main task, we collected the same additional measures as in study 1: the participants completed an extensive questionnaire about their second-order expectations (i.e. beliefs about the Investor's expectations) and fairness norms in the FBMTG; they completed a computerized version of the Social Value Orientation 'slider measure' task; and filled out the Guilt Inventory, the Dispositional Greed Scale, and a debriefing questionnaire. In this final step, they self-reported how they made their decisions in the main task by dividing 100 points over various candidate strategies (see study 1). The same measures as in study 1 were taken to ensure and check that the participants believed the Investors were real. Two participants indicated that they did not believe the Investors were real other participants; their data were excluded from analysis.

Computational modeling, software use, and ethical approval were identical to study 1.

Data and code availability

All data and code necessary to reproduce the results in this paper will be made available via Github upon publication.

Acknowledgements

This research was supported by European Research Council project 313454 (to A.S.). We thank Marc-Lluís Vives for helpful comments on the manuscript.

References

1. Axelrod, R. & Hamilton, W. D. The Evolution of Cooperation. *Science* (80-.). **211**, 1390–1396 (1981).
2. Nowak, M. A. Five rules for the evolution of cooperation. *Science* **314**, 1560–3 (2006).

3. Trivers, R. The evolution of reciprocal altruism. *Q. Rev. Biol.* **46**, 35–57 (1971).
4. Fehr, E. & Schmidt, K. M. A theory of fairness, competition, and cooperation. *Q. J. Econ.* **114**, 817–868 (1999).
5. Bolton, G. & Ockenfels, A. ERC: A theory of equity, reciprocity, and competition. *Am. Econ. Rev.* **90**, 166–193 (2000).
6. Tricomi, E., Rangel, A., Camerer, C. F. & O'Doherty, J. P. Neural evidence for inequality-averse social preferences. *Nature* **463**, 1089–91 (2010).
7. Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E. & Cohen, J. D. The neural basis of economic decision-making in the Ultimatum Game. *Science (80-.)*. **300**, 1755–8 (2003).
8. Battigalli, P. & Dufwenberg, M. Guilt in games. *Am. Econ. Rev.* **97**, 170–176 (2007).
9. Chang, L. J., Smith, A., Dufwenberg, M. & Sanfey, A. G. Triangulating the neural, psychological, and economic bases of guilt aversion. *Neuron* **70**, 560–72 (2011).
10. Charness, G. & Dufwenberg, M. Promises and partnership. *Econometrica* **74**, 1579–1601 (2006).
11. van Baar, J. M., Chang, L. J. & Sanfey, A. G. The computational and neural substrates of moral strategies in social decision-making. *Nat. Commun.* **10**, 1483 (2019).
12. Falk, A. & Kosfeld, M. The Hidden Costs of Control. *Am. Econ. Rev.* **96**, 1611–1630 (2006).
13. Bowles, S. Policies designed for self-interested citizens may undermine 'The moral sentiments': Evidence from economic experiments. *Science (80-.)*. **320**, 1605–1609 (2008).
14. Bowles, S. & Polania Reyes, S. Economic incentives and social preferences: a preference-based Lucas critique of public policy. *CESifo Work. Pap.* (2009).
15. Bowles, S. & Polanía-Reyes, S. Economic incentives and social preferences: Substitutes or complements? *J. Econ. Lit.* **50**, 368–425 (2012).
16. Hofstede, G. *Culture's consequences: Comparing values, behaviors, institutions and organizations across nations*. (Sage publications, 2001).
17. Haidt, J., Graham, J. & Joseph, C. Above and below left-right: Ideological narratives and moral foundations. *Psychol. Inq.* **20**, 110–119 (2009).
18. Haidt, J. *The Righteous Mind: Why Good People Are Divided By Politics and Religion*. (Pantheon Books, 2012).
19. Graham, J., Haidt, J. & Nosek, B. A. Liberals and Conservatives Rely on Different Sets of Moral Foundations. *J. Pers. Soc. Psychol.* **96**, 1029–1046 (2009).
20. Herrmann, B., Thoeni, C. & Gächter, S. Antisocial punishment across societies. *Science (80-.)*. **319**, 1362–1367 (2008).
21. Henrich, J. *et al.* In search of homo economicus: behavioral experiments in 15 small-scale societies. *Am. Econ. Rev.* **91**, (2001).
22. Henrich, J. *et al.* Costly Punishment Across Human Societies. *Science (80-.)*. **1767**, 1767–1771 (2007).
23. Berg, J., Dickhaut, J. & McCabe, K. Trust, Reciprocity, and Social-History. *Games Econ. Behav.* **10**, 122–142 (1995).

24. Johnson, N. D. & Mislin, A. A. Trust games: A meta-analysis. *J. Econ. Psychol.* **32**, 865–889 (2011).
25. Nihonsugi, T., Ihara, a. & Haruno, M. Selective Increase of Intention-Based Economic Decisions by Noninvasive Brain Stimulation to the Dorsolateral Prefrontal Cortex. *J. Neurosci.* **35**, 3412–3419 (2015).
26. Driessen, J., Van Baar, J., Sanfey, A., Glennon, J. & Brazil, I. Moral strategies and psychopathic traits. *Revis.*
27. Van Lange, P. A. M. The pursuit of joint outcomes and equality in outcomes: An integrative model of social value orientation. *J. Pers. Soc. Psychol.* **77**, 337–349 (1999).
28. Seuntjens, T. G., Zeelenberg, M., Van De Ven, N. & Breugelmans, S. M. Dispositional greed. *J. Pers. Soc. Psychol.* **108**, 917–933 (2015).
29. Jones, W. H., Schratter, A. K. & Kugler, K. The Guilt Inventory. *Psychol. Rep.* **87**, 1039–1012 (2000).
30. Murphy, R., Ackermann, K. J. & Handgraaf, M. J. Measuring social value orientation. *Judgm. Decis. Mak.* **6**, 771–781 (2011).
31. Van Lange, P. A. M., Otten, W., De Bruin, E. M. & Joireman, J. A. Development of prosocial, individualistic, and competitive orientations: theory and preliminary evidence. *J. Pers. Soc. Psychol.* **73**, 733–46 (1997).
32. Sears, D. O. & Funk, C. L. Evidence of the Long-Term Persistence of Adults' Political Predispositions. *J. Polit.* **61**, 1–28 (1999).
33. Galizzi, M. M. & Navarro-Martinez, D. On the External Validity of Social Preference Games: A Systematic Lab-Field Study. *Manage. Sci.* (2018). doi:10.1287/mnsc.2017.2908
34. Klein, O. *et al.* Low Hopes, High Expectations: Expectancy Effects and the Replicability of Behavioral Experiments. *Perspect. Psychol. Sci.* **7**, 572–584 (2012).
35. Zizzo, D. J. Experimenter demand effects in economic experiments. *Exp. Econ.* **13**, 75–98 (2010).
36. Baumeister, R. F., Stillwell, A. M. & Heatherton, T. F. Guilt: an interpersonal approach. *Psychol. Bull.* **115**, 243–67 (1994).
37. Gneezy, U., Meier, S. & Rey-Biel, P. When and Why Incentives (Don't) Work to Modify Behavior. *J. Econ. Perspect.* **25**, 191–210 (2011).
38. Gneezy, U. & Rustichini, A. A fine is a price. *J. Legal Stud.* **29**, 1 (2000).
39. Langner, O. *et al.* Presentation and validation of the radboud faces database. *Cogn. Emot.* **24**, 1377–1388 (2010).
40. Dufwenberg, M. & Gneezy, U. Measuring Beliefs in an Experimental Lost Wallet Game. *Games Econ. Behav.* **30**, 163–182 (2000).
41. Akaike, H. A New Look at the Statistical Model Identification. *IEEE Trans. Automat. Contr.* **19**, 716–723 (1974).
42. Hurvich, C. M. & Tsai, C.-L. Trust Regression and Time Series Model Selection in Small Samples. *Biometrika* **76**, 297–307 (1989).

43. Wagenmakers, E.-J. & Farrell, S. AIC model selection using Akaike weights. *Psychon. Bull. Rev.* **11**, 192–196 (2004).
44. Jones, E., Oliphant, T., Peterson, P. & Others. SciPy: Open source scientific tools for Python. (2001).