

Original citation:

Perera, Pri, Canic, Emina and Ludvig, Elliot Andrew, 1975-. (2015) *Cruel to be kind but not cruel for cash : harm aversion in the dictator game*. *Psychonomic Bulletin & Review* .

Permanent WRAP url:

<http://wrap.warwick.ac.uk/73453>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

"The final publication is available at Springer via <http://dx.doi.org/10.3758/s13423-015-0959-4> ."

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP url' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: publications@warwick.ac.uk

warwick**publications**wrap

highlight your research

<http://wrap.warwick.ac.uk>

Cruel to be kind but not cruel for cash: Harm aversion in the dictator game

Pri Perera, Emina Canic, Elliot A. Ludvig

Department of Psychology, University of Warwick

Word Count: 3996

Abstract

People regularly take prosocial actions, making individual sacrifices for the greater good. Similarly, people generally avoid causing harm to others. These twin desires to do good and avoid harm often align, but can sometimes diverge, creating situations of moral conflict. Here, we examined this moral conflict using a modified dictator game. Participants chose how much money to allocate away from a recipient, who was designated as an orphan, creating a sense of harm. This money was then re-allocated to either the participant or a charity. People were strongly prosocial: they allocated more money away from the orphan for charity than for themselves. Furthermore, people left more money with the orphan when the harm was framed as a means (taking) than as a side effect (splitting). As predicted by dual-process theories of moral decision making, response times were longer with the take action and positively correlated with the amount taken from the orphan. We conclude that just as people take positive actions for the greater good, they are similarly more willing to cause harm when it benefits others rather than themselves.

Keywords: Moral decision-making, harm aversion, prosocial behaviour, dictator game, dual-process models

Cruel to be kind but not cruel for cash: Harm aversion in the dictator game

People are generally averse to harming others, a tendency fostering the social cooperation upon which modern society is based (Greene, 2014; Nowak & Sigmund, 2005). But harm can arguably achieve a greater good, as evident in policies from taxation to military intervention. In such situations, the aversion to harm is placed in opposition with the desire to do good, leading to moral conflict. The Robin Hood legend famously exemplifies this conflict: Taking from the rich for the greater good of giving to the poor is lionized, whereas the parallel act of taking for ourselves would be condemned (Brickman & Bryan, 1975). Often though, people harm less directly, perhaps as a side effect of prioritising personal material desires over the basic needs of others (Singer, 1972).

Nonetheless, people do sometimes put others above themselves. They get happiness from spending money on others (Dunn et al., 2008) and work harder on menial tasks when earning for charity (Tonin & Vlassopoulos, 2013). People even cheat more in dice games when winnings go to charity (Lewis et al., 2012) and pay more to prevent shocks to others than themselves (Crockett et al., 2014). Using a novel variant of the dictator game, we show that people are also more willing to cause harm for the benefit of others than for themselves.

Theories of moral cognition posit a dual-system framework with separable contributions of actions and outcomes (Crockett, 2013; Cushman, 2013; Greene, 2014). Moral choice results from a compromise between a reflexive response to the permissibility of an action and a reflective calculation of the utilitarian costs/benefits incurred by its outcome. This framework derives support from studies showing that harm is judged less wrong when caused as a side effect than as a means, despite identical consequences (Greene, 2014). Take the classic trolley problem (Foot, 1967; Thompson, 1985): killing one person to save five by diverting a runaway trolley onto someone is judged more morally permissible than pushing someone onto the tracks to block the trolley (Greene et al., 2009). Further evidence for this

framework comes from studies showing that increased deliberation time leads to more utilitarian judgments (Rand et al., 2012; Suter & Hertwig, 2011). Many of these studies, however, use hypothetical dilemmas, which lack realism and struggle to replicate actual consequential choices (Bauman, et al., 2014; Gold et al., 2014; Kang & Camerer, 2013).

Social economic games provide an alternative approach for evaluating moral choices. They are advantageous because they involve real-world, monetary consequences and do not enforce all-or-nothing resolutions. For example, in the dictator game (Kahneman et al., 1986), a participant is given money to distribute between themselves and a passive recipient. People generally give something, but usually less than an equal split (Engel, 2011). Harm can be evoked in these games by initially designating the money involved to the recipient, thus changing the distribution action from giving to taking (Keysar et al., 2008; Leliveld et al., 2008).

Such framing is usually consequential in related economic games (e.g. Andreoni, 1995; Cookson, 2000; Leliveld et al., 2008; although not always, e.g., Rubinstein, 2014). In the dictator game, however, mixed results have emerged. Some studies have found that including a take action in the dictator game caused dictators to behave more selfishly (Bardsely, 2008; Krupka & Weber, 2013; List, 2007), but these provided both give and take options simultaneously, making selfish choices appear comparatively less harmful (Parducci, 1965). When choice options have been kept consistent, simply framing the game as either giving or taking had no effect on distributions (Chowdhury et al., 2014; Dreber et al., 2013; Grossman & Eckel, 2015; Smith, 2015).

This study adapted the standard dictator game in three ways. First, the recipient was designated as an orphan needing charitable help, making the choice more consequential and increasing the sense of harm elicited from any money not allocated to her. Second, a condition was included where participants allocated money between the orphan and a charity.

Donations to the charity were described as benefiting more orphans than equivalent donations made to the individual orphan, retaining the tension between an aversive action (taking from an orphan) and a greater good (giving to charity) that characterizes sacrificial moral dilemmas (Foot, 1967). Third, to distinguish harm caused as a means to an end and harm caused as a side effect, two conditions cast the distribution action involved in the game differently: either as taking or splitting. In the take action, the orphan was initially provided money and the participant decided how much to take away, directly causing harm as a means to an end. In the split action, initial ownership was neutral, making any non-donation to the orphan a harmful side effect of the distribution decision.

The experiment directly compares the strengths of self-interest and prosociality in mitigating harm. If prosociality outweighs self-interest then, succeeding Robin Hood, people should allocate less money to the orphan (be less harm averse) when distributing money to charity than when distributing money to themselves. By the dual-process framework, people should be more conflicted in the take condition, where the harm serves as means to an end. Thus, they should leave more money with the orphan in the take conditions. Moreover, if people are indeed overcoming an intuitive aversion to directly causing harm, the less money they allocate to the orphan, the longer they should take to respond. Finally, if prosocial motives are specific to mitigating harm caused as a means to an end, any extra amount people allocate to charity over themselves should be larger in the take condition than the split condition.

Method

Participants.

800 participants were recruited using Amazon Mechanical Turk. This number was determined in advance due to a fixed budget for data collection. 124 total participants were excluded for not meeting the exclusion criteria, which were determined in advance: duplicate

IP addresses (52 participants), failed attention (67) or transcription test (39), or answered the main task too quickly (<8 seconds based on a separate pilot study) to have read the questions (10). Including these participants did not substantively change the results. Of the 676 participants (410 men; 263 women; 3 undisclosed) retained for analyses, ages ranged from 18 to 65+ (mean = 26-30 years). 652 were born, had lived longest and were currently living in the US. Participants were paid \$0.50 for the 5-minute task and could earn up to \$1 extra either for themselves or charity, depending on the condition. Informed consent was obtained, and the study was approved by the appropriate University ethics committee.

Procedure.

A 2-x-2 between-subjects design was used that varied the incentive (self vs. charity) and the distribution action (split vs. take) in a one-shot dictator game. In the *self* conditions, participants were asked to distribute \$1 between themselves and an orphan. The *charity* conditions were identical to the self conditions, except that participants instead proposed a distribution between the orphan and a charity. The distributive action was described using the stated initial ownership of the \$1. In the split conditions, initial ownership was neutral. In the take conditions, initial ownership was with the orphan. The primary dependent variable was the amount remaining with the orphan once participants had completed the task.

After providing consent, participants were randomly assigned to one of four groups: split-self ($n=178$), take-self ($n=172$), split-charity ($n=171$), take-charity ($n=155$). Participants were reminded that all money involved was real and would be distributed according to their decision.

Participants were then asked to distribute 100 cents with the exact wording varying by condition (see Supplementary Materials). Below the question was a photo of an orphan along with her name, a brief biography, and a statement explaining that 100 cents would provide four meals. Beneath this section, in the *self* condition, the word “you” was displayed with a

generic avatar. In the *charity* conditions, a charity logo, name, and a brief description that explained they could give 10 children a meal with 100 cents was displayed. Adjacent to each of the two donation options was a textbox where participants could type how much money they wished to donate. There was also a *Total* textbox that automatically displayed the sum of the two donations. Participants were only able to progress if this total displayed 100. The order in which each entity appeared was randomised across participants.

For the split conditions, both textboxes were initially set to 0 cents. Participants were told that they had been given 100 cents and could divide it between themselves/charity and the orphan as they pleased. They were asked to indicate how they would like to allocate the endowment by typing in the relevant textboxes. The take condition differed in that the orphan's textbox was initially set to 100 cents while the participant's/charity's was set to 0 cents. Participants were told that the orphan had been given a donation of 100 cents and that they could take any amount of this away from her for themselves/charity. All response times were recorded, and outliers (>3 SD from the mean, 10 in total) were removed from the response-time analyses.

Following the main task, participants were given a basic arithmetic test to filter for task comprehension and an arbitrary effort task to filter out automated responses and reckless participants (as in Rand et al., 2012). Participants were also asked what the average distribution behaviour of others completing the same task would be and given a demographics questionnaire. Allocations made to themselves were duly received. Allocations made to charity or the orphan were donated to a children's charity.

Results are reported as the mean \pm 95% confidence intervals. Inferential statistics for null-hypothesis significance testing are reported for completeness, for which an alpha level of .05 was taken as statistically significant. A Bonferroni correction was applied to adjust for

multiple comparisons when needed. Procedural details and initial analyses were decided in advance and recorded on the Open Science Framework (<https://osf.io/z7nfd/>).

Results

Figure 1 shows that with both the split and take distribution actions, people allocated more money for charity than for themselves. The amount left with the orphan increased from 42.0 ± 3.7 cents in the charity conditions to 53.6 ± 4.6 cents in the self conditions. A two-way ANOVA confirmed that people were reliably more willing to cause harm when the charity was the beneficiary than when they themselves were, independent of the distribution action, $F(1, 672) = 14.28, p < .001, \eta_p^2 = .021$.

In addition, changing the distribution action (take vs. split) shifted the amount of money remaining with the orphan. Overall, people left less money with the orphan in the split condition (40.7 ± 3.9 cents) than in the take condition (55.8 ± 4.6 cents), $F(1, 672) = 25.38, p < .001, \eta_p^2 = .036$. The same general pattern was also observed when people were asked what others would do (see Figure S1). People were less willing to cause harm when faced with taking money from the orphan than when faced with splitting the money.

There was no reliable interaction between the incentive and the distribution action, $F(1, 672) = 0.15, p = .70, \eta_p^2 = .0002$. The extra amount allocated to charity (over self) was not different based on the distribution action ($MD_{Take} = 12.6 \pm 9.0$ cents; $MD_{Split} = 10.3 \pm 7.7$ cents). Thus, the harm mitigating effects of prosocial motivations did not depend on whether the harm was construed as a means to an end or a side effect.

Figure 2 shows that people took longer to respond with the take action and in the charity conditions. Response times rose from 33.3 ± 2.2 s with the split action to 38.5 ± 2.8 s with the take action, $F(1, 662) = 10.87, p < .05, \eta_p^2 = 0.01$, and decreased from 43.7 ± 2.6 s in the charity conditions to 28.5 ± 2.1 s in the self conditions, $F(1, 662) = 83.47, p < .05, \eta_p^2 = 0.11$. Furthermore, there was an interaction between incentive and action on response times,

$F(1,662)=4.89, p<.05, \eta_p^2=0.01$. In the charity conditions, response times were longer for the take than the split action ($MD_{Charity} = 9.3\pm 5.2$ s), but not in the self conditions ($MD_{Self} = 2.0\pm 4.1$ s).

These response times correlated with the amount allocated to the orphan in both take conditions. Figure 3 displays a scatterplot of response times against amount allocated to the orphan for both incentives and actions; Table 1 details the exact correlations. In the take conditions, there were significant negative correlations between response times and the amount left with the orphan, indicating that those responding faster allocated more money to the orphan regardless of incentive. In the split conditions, neither correlation was significant and they trended in opposite directions.

As the take condition involved a potential default option (100 cents to the orphan) that may have contributed to the negative correlations, response-time analyses were repeated excluding all 127 (self=91; charity=36) allocations that were exactly equal to 100 cents. These correlations trended in the same direction, but were smaller (take-charity: $r(113)=-.21, p=.025$; take-self: $r(76)=-.05, p=.67$). This pattern suggests that the negative correlations were driven in part by those participants who allocated the full 100 cents to the orphan. To further evaluate whether these 100-cent allocations were indeed default responses, response times were compared between those who allocated 100 cents to the orphan in the take (127) and split (66) conditions. These response times were not reliably different ($M_{take}=30.3\pm 3.7$ s; $M_{split}=29.8\pm 5.7$ s; $t(191)=0.17, p>.05$), making it unlikely that participants in the take group were treating the 100-cent allocation as an explicit default option.

Discussion

There were two striking results in this novel variant of the dictator game. First, people allocated more money away from the orphan and to the charity than to themselves. This can be considered a generalised “Robin Hood effect” (Brickman & Bryan, 1975), whereby people

are more willing to cause harm when striving for a greater good. Secondly, casting the dictator game as taking as opposed to splitting resulted in more money being left with the orphan. This finding contrasts with previous studies on dictator game framing which have not found effects of pure action framing on the observed distributions (e.g., Chowdhury et al., 2014; Dreber et al., 2013), perhaps due to the reduced moral distance invoked by having an orphan serve as recipient instead of an anonymous player (Aguiar et al., 2008). The sensitivity to the distribution action implies that people were more averse to causing harm as a means to an end (taking) than as a side effect (splitting). Consistent with studies using hypothetical life-and-death dilemmas such as the trolley problem, people's judgments regarding harm aversion were also sensitive to this distinction for economic incentives and consequences less severe than death (Gold, Pulford & Colman, 2013).

The interaction between incentive and action in response times and their correlations with the orphan allocations support a dual-system moral framework (Crockett, 2013; Cushman, 2013). A core idea of this theory is that moral judgments are based on intuitive, reflexive responses that can be overridden by a more utilitarian, reflective process (e.g. Greene, 2014), which takes deliberation time (Rand et al., 2014; Suter & Hertwig, 2011). Here, the impulse to avoid harm (taking from the orphan) could be overridden by concern for the greater good (giving to charity).

Participants took longer to respond in the take condition than in the split condition, but only in the charity conditions (see Figure 2). This pattern suggests that participants took longer to resolve the conflict between a directly harmful action and the greater good. More strikingly, response times in the take conditions were negatively correlated with the amount left with the orphan for both the self-interested and prosocial incentive. Overcoming the intuitive aversion to directly taking from an orphan took time: responding quicker left people more harm averse and thus less likely to benefit the greater good. Splitting, which only

caused harm as a side effect, presented no reliable correlations. These observations provide evidence for a dual-process framework in a new domain with consequential economic harm (Rand et al., 2014).

Contrary to our initial hypothesis, there was no interaction in the amount allocated between the incentive and the distribution action. We had expected the prosocial motivation to more strongly mitigate harm produced as a means than as a side effect. Instead, the prosocial motivation mitigated both types of harms equally.

A limitation of the current design is that the take conditions had a potential default response. To emphasize harm caused by taking, the orphan was initially endowed 100 cents, whereas in the split conditions their initial endowment was 0. Thus, in the take conditions, allocating any money from the orphan involved overriding a default. The observed effect of distribution action could therefore in part represent the participants' adherence to this default (Dhingra et al., 2012). Indeed, more participants allocated the full 100 cents to the orphan in the take than in the split conditions.

Two factors mitigate this interpretation. First, adherence to a default should yield shorter response times, yet participants took longer to respond in the take condition (Figure 2). Even amongst those who allocated the full 100 cents, there was no difference in response times between the take and split conditions. Second, when removing all subjects who allocated the full 100 cents to the orphan, the negative correlation between allocations and response-times holds in the take-self condition. The negative correlation in the take-charity condition becomes non-significant, however. This suggests that the correlation in the full group (Table 1) was driven in part by quicker responses from those who made the maximum allocation to the orphan. Thus, a default option interpretation for this subset of the results cannot be entirely ruled out. Indeed, one possibility is that violating the default norm is part of what makes taking an aversive action (Baron & Ritov, 2004; Mazar & Hawkins, 2015).

Allocations to charity were specified as benefitting the greater good in that they provided more meals to more orphans. This more efficient distribution of the available resources was clearly the rational, utilitarian distribution, yet people were still hesitant to allocate money away from the orphan. Nevertheless, they were more willing to cause harm by allocating money to this prosocial cause than they were for selfish reasons. Thus, just as people can be more motivated to do positive deeds for others than for themselves (Dunn et al., 2008; Tonin & Vlassopoulos, 2013), they can be similarly more motivated to cause harm to benefit others.

Provided context, people are more willing to cause harm to benefit the greater good than themselves. Intriguingly, the prosocial motivation in this study mitigated equally against harm caused as a means to an end and as a side effect. These findings suggest that, despite traditional economics encouraging policy designs incentivising self-interest (Bowles, 2008), creating psychological barriers to self-serving behaviour and emphasizing prosocial values may be more effective.

References

- Andreoni, J. (1995). Warm-glow versus cold-prickle: The effects of positive and negative framing on cooperation in experiments. *Quarterly Journal of Economics*, *110*, 1-21.
- Aguiar, F., Brañas-Garza, P., & Miller, L. (2008). Moral distance in dictator games. *Judgement and Decision Making*, *3*(4), 344-354.
- Bardsley, N. (2008). Dictator game giving: altruism or artefact? *Experimental Economics*, *11*(2), 122-133.
- Baron, J., & Ritov, I. (2004). Omission bias, individual differences, and normality. *Organizational Behavior and Human Decision Processes*, *94*, 74-85.
- Bauman, C. W., McGraw, A. P., Bartels, D. M., & Warren, C. (2014). Revisiting external validity: Concerns about trolley problems and other sacrificial dilemmas in moral

- psychology. *Personality and Social Psychology Compass*, Forthcoming. Available at SSRN: <http://ssrn.com/abstract=2478465>
- Bowles, S. (2008). Policies designed for self-interested citizens may undermine “the moral sentiments”: Evidence from economic experiments. *Science*, *320*(5883), 1605-1609.
- Brickman, P., & Bryan, J. H. (1975). Moral judgement of theft, charity, and third-party transfers that increase or decrease equity. *Journal of Personality and Social Psychology*, *31*(1), 156-161.
- Chowdhury, S. M., Jeon, J. Y., & Saha, B. (2014). Giving, taking and gender in dictator games. CBESS Discussion Paper, 14-09. Available at: <https://www.uea.ac.uk/documents/166500/0/CBESS-14-09/753b730e-599d-4462-9ebd-201357ce6f79>
- Cookson, R. (2000). Framing effects in public goods experiments, *Experimental Economics*, *3*, 55-79.
- Crockett, M. J. (2013). Models of morality. *Trends in Cognitive Science*, *17*(8), 363-366.
- Crockett, M. J., Kurth-Nelson, Z., Siegel, J. Z., Dayan, P., & Dolan, R. (2014). Harm to others outweighs harm to self in moral decision-making. *Proceedings of the National Academy of Sciences, USA*, *111*, 17320-17325.
- Cushman, F. A. (2013). Action, outcome, and value: A dual-system framework for morality. *Personality and Social Psychology Review*, *17*(3), 273-292.
- Dhingra, N., Gorn, Z., Kener, A., & Dana, J. (2012). The default pull: An experimental demonstration of subtle default effects on preferences. *Judgement and Decision Making*, *7*(1), 69-76.
- Dreber, A., Ellingsen, T., Johannesson, M., & Rand, D. (2013). Do people care about social context? Framing effects in dictator games. *Experimental Economics*, *16*(3), 349-371.

- Dunn, E. W., Aknin, L. B., & Norton, M. I. (2008). Spending money on others promotes happiness. *Science*, *319*, 1687-1688.
- Engel, C. (2011). Dictator games: A meta study. *Experimental Economics*, *14*, 583-610.
- Foot, P. (1967). The problem of abortion and the doctrine of the double effect in virtues and vices. *Oxford Review*, *5*, 5-15.
- Gold, N., Pulford, B. D., & Colman, A. M. (2014). The outlandish, the realistic, and the real: Contextual manipulation and agent role effects in trolley problems. *Frontiers in Psychology: Cognitive Science*, *5*(35), 1-10.
- Gold, N., Pulford, B. D., & Colman, A. M. (2013). Your money or your life: Comparing judgements in trolley problems involving economic and emotional harms, injury and death. *Economics and Philosophy*, *29*, 213-233.
- Greene, J. D. (2014). *Moral tribes: Emotion, reason, and the gap between us and them*. London: Atlantic Books Ltd.
- Greene, J. D., Cushman, F. A., Stewart, L. E., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2009). Pushing moral buttons: The interaction between personal force and intention in moral judgment. *Cognition*, *111*(3), 364-371.
- Greene, J. D., Morelli, S. A., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2008). Cognitive load selectively interferes with utilitarian moral judgment. *Cognition*, *107*(3), 1144-1154.
- Grossman, P. J., & Eckel, C. C. (2015). Giving versus taking for a cause. *Economics Letters*, *132*, 28-30.
- Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1986). Fairness and the assumptions of economics. *Journal of Business*, *59*, S285-S300.
- Kang, M. J., & Camerer, C. F. (2013). fMRI evidence of hot-cold empathy gap in hypothetical and real choices. *Frontiers in Neuroscience*, *7*(104), 1-16.

- Keysar, B., Converse, B. A., Wang, J., & Epley, N. (2008). Reciprocity is not give or take: Asymmetric reciprocity to positive and negative acts. *Psychological Science, 19*(12), 1280-1286.
- 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*(3), 495-524.
- Leliveld, M. C., van Dijk, E., & van Beest, I. (2008). Initial ownership in bargaining: Introducing the giving, splitting, and taking ultimatum bargaining game. *Personality and Social Psychology Bulletin, 38*(9), 1214-1225.
- Lewis, A., Bardis, A., Flint, C., Mason, C., Smith, N., Tickle, C., & Zinser, J. (2012). Drawing the line somewhere: An experimental study of moral compromise. *Journal of Economic Psychology, 33*(4), 718-725.
- List, J. A. (2007). On the interpretation of giving in dictator games. *Journal of Political Economy, 115*(3), 482-493.
- Mazar, N., & Hawkins, S. A. (2015). Choice architecture in conflicts of interest: Defaults as physical and psychological barriers to (dis)honesty. *Journal of Experimental Social Psychology, 59*, 113-117.
- Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature, 437*, 1291-1298.
- Parducci, A. (1965). Category judgment: A range-frequency model. *Psychological Review, 72*(6), 407-418.
- Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature, 489*, 427-430.

- Rand, D. G., Peysakhovich, A., Kraft-Todd, G. T., Newman, G. E., Wurzbacher, O., Nowak, M. A., & Greene, J. D. (2014). Social heuristics shape intuitive cooperation. *Nature Communications*, 5(3677), 1-12.
- Rubinstein, A. (2014). A typology of players: Between instinctive and contemplative. Available at: www.arielrubinstein.tau.ac.il/papers/Typology.pdf
- Singer, P. (1972). Famine, affluence, and morality. *Philosophy and Public Affairs*, 1(3), 229-243.
- Smith, A. (2015). On the nature of pessimism in taking and giving games. *Journal of Behavioral and Experimental Economics*, 54, 50-57.
- Suter, R. S., & Hertwig, R. (2011). Time and moral judgement. *Cognition*, 119, 454-458.
- Thomson, J. J. (1985). The trolley problem. *Yale Law Journal*, 94, 1395-1415.
- Tonin, M., & Vlassopoulos, M. (2013). Social incentives matter: Evidence from an online real effort experiment. *Review of Environment, Energy and Economics*, Re3, Forthcoming. Available at SSRN: <http://ssrn.com/abstract=2209865>

Author Contributions

The initial idea was conceived by PP and was honed into the finalised design through several fruitful group discussions with EAL and EC. The experimental protocol was coded by PP, who then collected the data and analysed it. EC followed up with further analysis and developed figures to present the results. Finally, PP drafted the manuscript, which was subsequently revised by EAL and EC, who iteratively provided significant contributions and extensive, critical revisions. All authors contributed to re-working the revision.

Table 1. *Correlations between response times and amount left with the orphan in the different experimental conditions.*

	Take	Split
Self	-.26**	.16
Charity	-.20*	-.05

Note. * $p < .05$; ** $p < .001$, after Bonferroni correction.

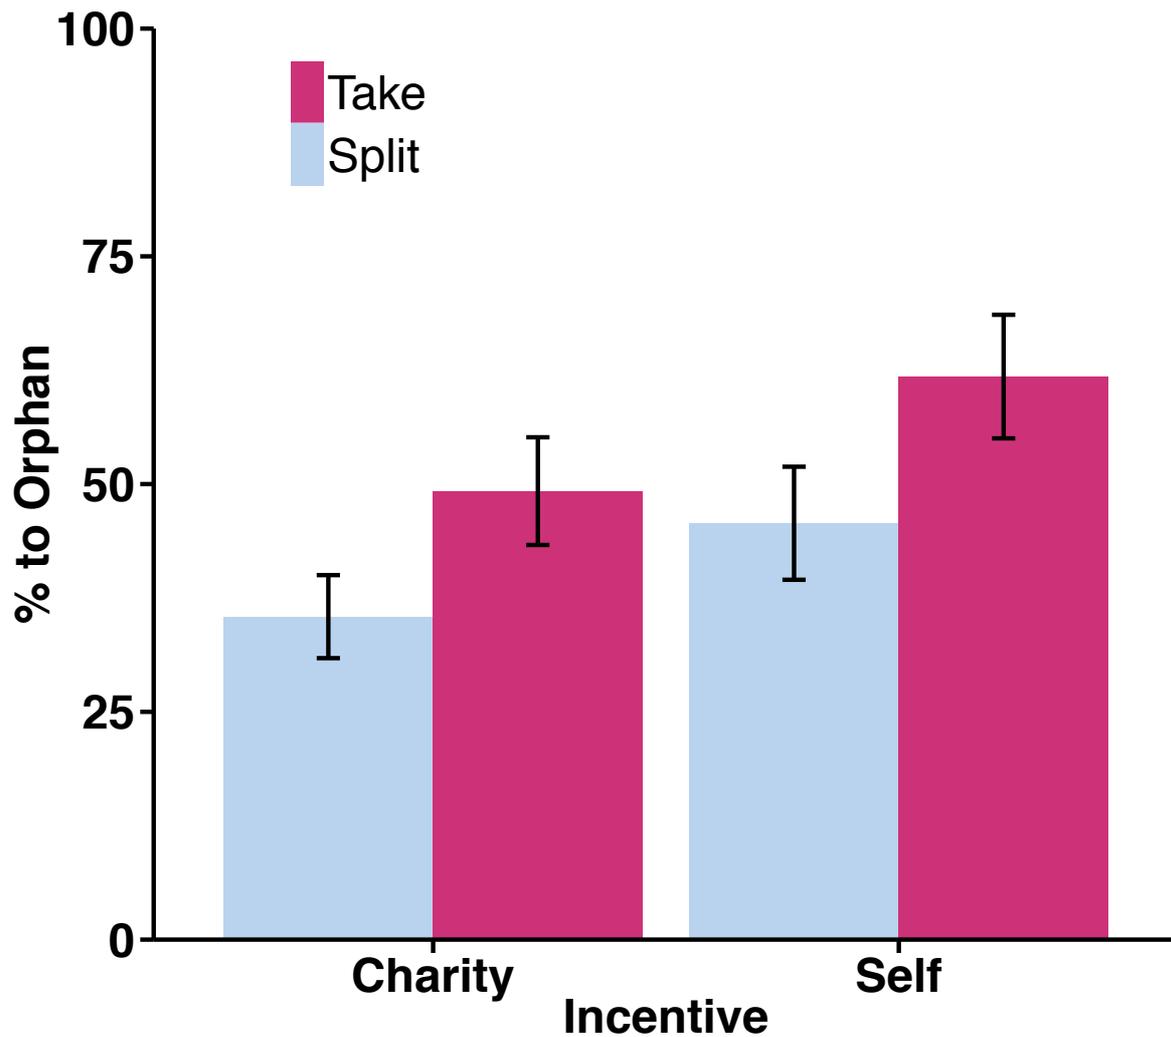


Figure 1. Mean amount of money left with the orphan (\pm 95% CI) as a function of incentive (self vs. charity) and action (take vs. split). The orphan received more money when allocations were made between her and the participant than when allocations were made between her and the charity. Participants allocated more money to the orphan in the take conditions than in the split conditions.

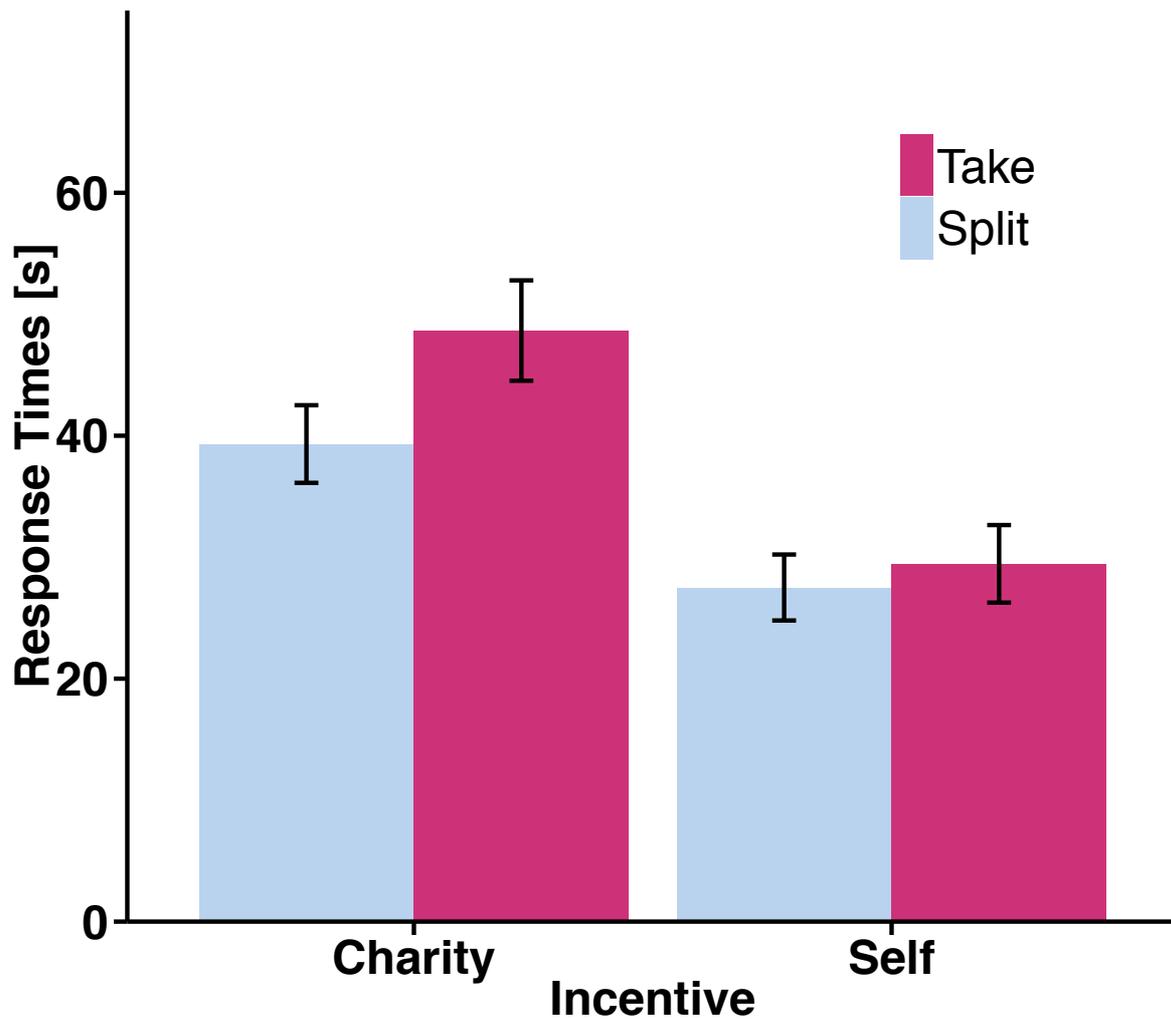


Figure 2. Mean response times (\pm 95% CI) as a function of incentive (self vs. charity) and action (take vs. split). When allocating money between charity and the orphan, people took longer to respond with the take action. When allocating money between themselves and the orphan, however, response times did not reliably differ based on the distributive action.

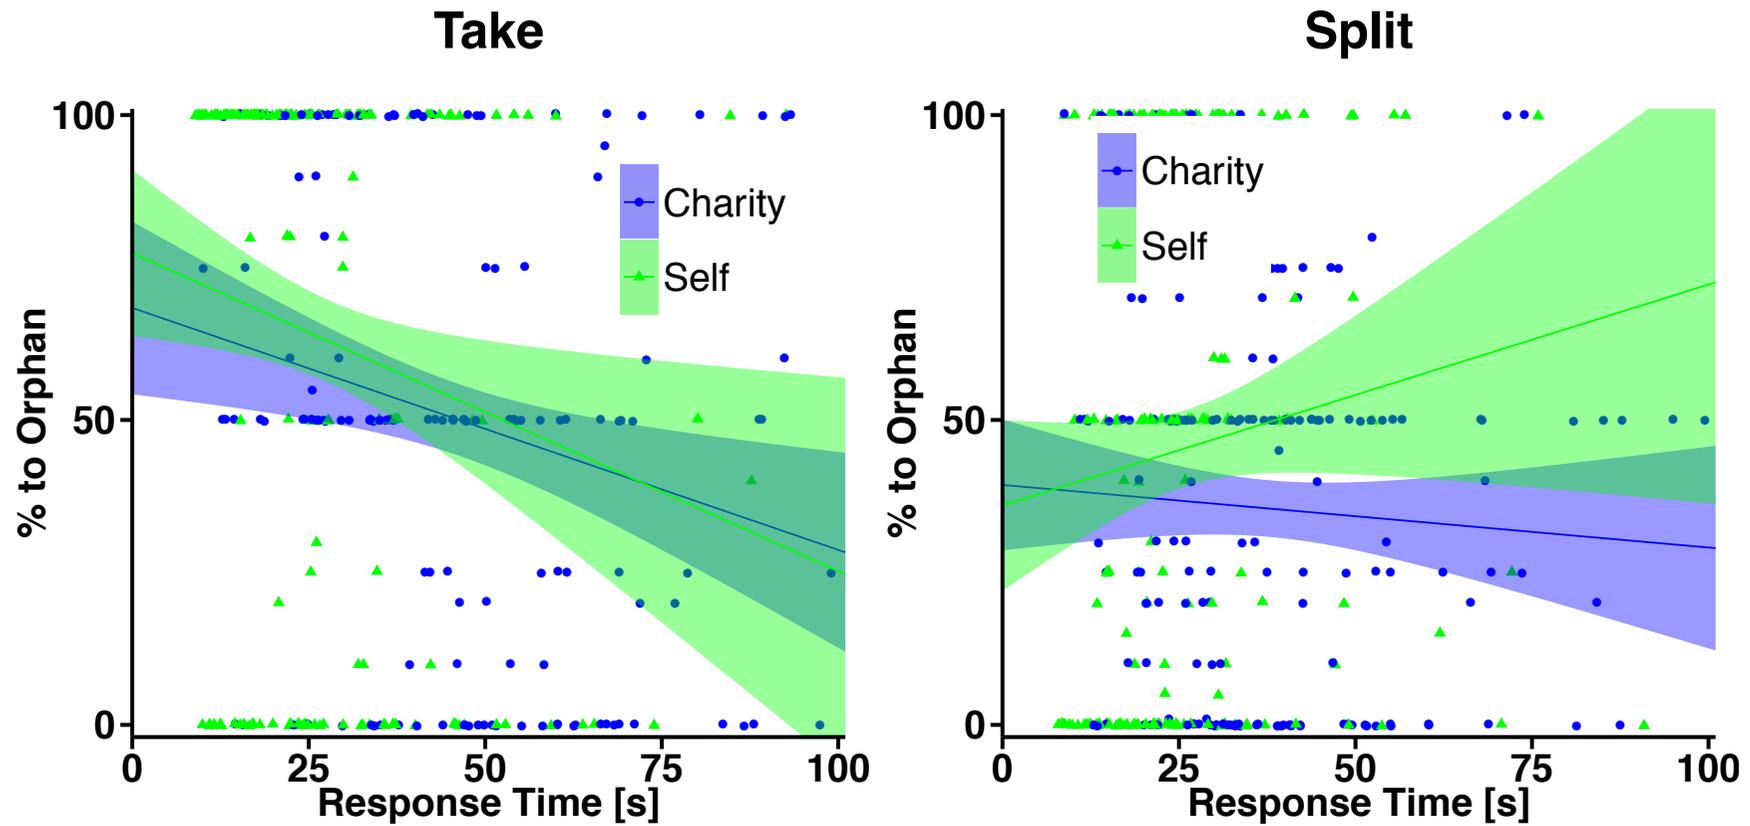


Figure 3. Scatterplot of amount left with the orphan as a function of incentive (charity or self) and response times. The left panel is for the split action, and the right panel is for the take action. Lines are the best-fitting regression lines with the shaded areas representing 95% confidence intervals. Note that many people chose to allocate all, none, or a 50/50 split, hence the cluster of dots along those three lines in both plots.