

# Lying Behavior When Payoffs are Shared with Charity: Experimental Evidence

Scott Lee Chua<sup>a</sup>, Jessica Chang<sup>a</sup>, and Guillem Riambau<sup>\*b,c</sup>

<sup>a</sup>Yale-NUS College

<sup>b</sup>Universitat de Barcelona

<sup>c</sup>Institutions and Political Economy Research Group

December 9, 2020

## Abstract

We investigate lying behavior when lying is undetectable and payoffs are split with charity. 524 participants roll a die in private, report the outcome, and receive the monetary equivalent of their reported number, i.e., there is a clear incentive to lie. Participants are randomly assigned to share all, some, or none of this payoff with a charity of their choice. This allows us to examine how lying behavior changes with the share of payoffs going to charity. Our results are as follows: (i) there are participants in every group who lie to inflate their reported number; (ii) participants with no share of the payoff lie much less than participants with some share, no matter how small; and (iii) post-experiment surveys reveal that participants who keep the whole payoff are much less likely to admit to having cheated than all other participants. Finally, our data suggests that lying is not correlated with any observable sociodemographic characteristic.

**Keywords:** Dishonesty, Charity, Morality, Split, Experiment

**JEL Codes:** C91, D63, D64

**Authors' declarations of interest:** None.

**Funding:** We gratefully acknowledge funding from Yale-NUS College. The funding source had no involvement in the design, collection, analysis, interpretation, writing, or decision to submit this article.

**Data and code availability:** Anonymized data and replication codes are available as online resources.

---

\*Corresponding author. Email: griambau@gmail.com. ORCID: 0000-0002-4076-0616. Present affiliation: Universitat de Barcelona, Gran Via de les Corts Catalanes, 585, Barcelona 08007. The authors would like to thank Eugene Choo, Lam Wingtung, and Sam Guenther for their helpful insights; and Minwoo Choi, Dylan Leong, and Yogesh Tulsi for their assistance in conducting the experiment. All errors remain our own.

# 1 Introduction

In traditional economic models, people are predicted to lie so long as the lie goes undetected and there are material gains to be made. Recent studies have shown, however, that people actually lie much less than expected: a meta-analysis of 90 experiments involving 44,390 subjects across nearly 50 countries shows that subjects forgo on average three-quarters of the potential private gains from lying, even when lying is both incentivized and undetectable (Abeler et al., 2019). Although the motives behind lying aversion are yet to be fully understood, the general consensus in the literature is that people refuse to lie not only because of an intrinsic preference for being honest, but also because they want to be perceived as honest by others (Abeler et al., 2019; Gneezy et al., 2018).

Another strand of the literature has sought to understand “other-regarding preferences,” i.e., how one’s choices are affected by the welfare of others. Evidence from over one hundred dictator games consisting of over 20,000 observations reveals that over two-thirds of dictators give a strictly positive amount, with an (unconditional) average donation of slightly less than one third of the pie<sup>1</sup> (Engel, 2011). This result also contradicts standard economic models that assume people only regard private material benefits.

This paper unites these two strands of the literature by examining lying behavior when, by means of undetectable cheating, people can benefit others beyond themselves. Recent studies have begun to compare *selfish* lying, where payoffs are kept privately in full, with *pure prosocial* lying, where payoffs are donated to other anonymous individuals. Wiltermuth (2011)<sup>2</sup> finds higher rates of selfish lying than pure prosocial lying, while Gino et al. (2013)<sup>3</sup> find similar rates for both. However, the same studies<sup>4</sup> find the highest rate of lying when payoffs are split evenly between the participant and another individual – up to 50% more frequent than selfish lying. These results suggest that lying solely for oneself is more likely than lying solely for another, but less likely than when payoffs are split.

Our study examines lying behavior when the other recipient is a *charity* rather than an individual. “Since private philanthropy can substitute for public sector provision of goods and services,” Andreoni (2006) notes, “it becomes essential to understand how private charity is provided” (p.1205). We design our experiment following Konow (2010), who argues that

---

<sup>1</sup>Levitt and List (2007) report a similar pattern.

<sup>2</sup>Study 2.

<sup>3</sup>Experiment 3.

<sup>4</sup>Wiltermuth (2011) Studies 2 and 3; Gino et al. (2013) Experiments 2 and 3.

involving real charities brings lab experiments closer to the real world.

Similar to Klein et al. (2017), we implement a range of split-payoff treatments to observe how marginal variations in payoff splits might affect lying behavior. Each participant in our experiment is randomly assigned to one of five groups. In the first group, participants keep all payoffs. In the second group, participants donate all payoffs to charities of their choice. In the three remaining groups, payoffs are split between participants and their chosen charities: 90%–10%, 50%–50%, and 10%–90%, respectively.

We implement a die-rolling task following the Fischbacher and Föllmi-Heusi (2013) paradigm. In this task, participants privately roll a die and report the outcome, the monetary equivalent of which they either keep, donate, or split with charity, depending on their treatment group. Participants then receive their share of the money, and are requested to oversee the donation to charity, if any. When running the experiment, we made every effort to ensure that participants (i) were aware that lying was possible and completely undetectable; (ii) would care about the charity with whom they would share their payoffs; and (iii) trusted that we would indeed make the donation as specified. Although we cannot detect dishonesty on an individual level, we can measure dishonesty at the aggregate level by comparing the distribution of reported outcomes against the expected discrete uniform distribution of a fair die roll.

We find that participants in all treatment groups unambiguously lie by inflating their reported outcomes. However, lying rates are heterogeneous with respect to payoff split. We estimate that when participants privately benefit from the lie (whether partially or in full), one in four participants who have an incentive to lie do so. However, only one in ten are estimated to lie when payoffs go entirely to charity. These results suggest a clear discontinuity in lying rates, with a sharp decrease when the participant no longer benefits from the lie.

Our results suggest that prosocial lying behavior might depend on the nature of the recipient. Lupoli et al. (2017) and Maggian (2019) also use a charity as the “other” recipient that stand to benefit from participants’ lies. Both studies find that participants lie at similar rates whether the lie benefits the charity or the self, suggesting that prosocial lies are no easier to internalize than selfish lies. To account for this, Maggian (2019) suggests that the greater psychological distance to organizations may hinder willingness to incur lying costs. In other words, participants may find it equally difficult to internalize a lie when the beneficiary is a faceless organization, no matter how noble. Our results support this finding, as we find that prosocial lying occurs less frequently than selfish lying.

Results in this paper also suggest that prosocial lying may depend on the nature of the task at hand. Gino et al. (2013) and Wiltermuth (2011) ask participants to report their task-solving scores (e.g., solving matrices), and in doing so, signal something about their effort or ability. Both studies find significantly higher lying rates when payoffs are evenly split, than when participants take all. In contrast, the tasks in Klein et al. (2017) and this study (coin flip and die roll, respectively) are luck-based and costless: we both find no increase in lying when payoffs are evenly split.

Finally, we find a kink: lying decreases sharply when payoffs go from 10% to self to 0% to self —i.e., entirely to charity. This result, in line with Klein et al. (2017), suggests that the extensive margin (whether participants receive *any* benefit from lying) plays a larger role than the intensive margin (the size of the benefit they receive) in participants' decision to lie.

By means of a post-experiment questionnaire, we also examine who is most likely to admit having lied. We estimate that one in four participants who lied admit to having done so, with the exception of those who take home all the payoff from lying. Among them, only one in thirty who lied admit to it. This suggests that when charity is involved, being honest and being seen as honest operate under different mechanisms. Participants who donate all payoff to charity are least likely to lie, while those who keep all payoff for themselves are least likely to admit to lying.

Our study thus contributes to the emerging experimental literature that focuses on the interplay between (dis)honesty and other-regarding preferences. Overall, our results support earlier findings that preferences to be seen as honest affect participants' lying behavior. More precisely, the identity of the receiver, the difficulty of the task, and the expectation that observers (e.g., experimenters) update their beliefs based on participants' reported outcomes are likely determinants of behavior in honesty games with external beneficiaries.

## 2 Experiment design

### 2.1 Treatment groups

Participants privately roll a fair six-sided die and report the outcome. The payoffs disbursed are the monetary equivalents of the outcomes reported (e.g., if a participant reports a 4, the corresponding payoff is \$4). Participants thus have opportunity and incentive to lie by misreporting the outcome. We vary whether or not a payoff is split between a participant and a

charity (and if so, in what proportions), and observe how lying behavior changes accordingly.

We randomly distribute participants across five different groups. In our control group (the “self-only” group), participants take home 100% of the payoff, consistent with previous experiments that use the die-roll paradigm (e.g., Hermann and Mußhoff, 2019). Participants in the three split-payoff groups (“90-self”, “50-self”, and “10-self”) take home only a fraction of the payoff (90%, 50%, and 10%, respectively), and donate the rest to a charity of their choice. In the last group (the “charity-only” group), participants are required to donate all their payoff, leaving them with no take-home share. Participants are aware of the payoff scheme they face before they roll the die.

Participants can choose from five reputable charities which represent a diverse range of social causes our target population (undergraduate university students) might care about: women’s rights, prisoner rehabilitation, animal welfare, crisis relief, and terminally ill children (Appendix A.5). All five charities accept online donations, which allows us to disburse donations in front of the participants in real-time. Table 1 summarizes the payoff schemes of the different treatment groups.

| Treatment Group        | Participant<br>(% payoff) | Charity<br>(% payoff) |
|------------------------|---------------------------|-----------------------|
| self-only              | 100                       | 0                     |
| split-payoff (90-self) | 90                        | 10                    |
| split-payoff (50-self) | 50                        | 50                    |
| split-payoff (10-self) | 10                        | 90                    |
| charity-only           | 0                         | 100                   |

Table 1: Comparison of payoff schemes in different treatment groups

## 2.2 Die-rolling task

Participants are brought inside a private room, one at a time, by an experimenter. Three items are provided inside the room: a pouch containing a die, an abridged copy of the task instructions (Appendix A.2), and a timed lock-box containing pen and paper. As the lock-box is transparent, participants can see the pen and blank slips of paper inside. The experimenter then sets the timer in the lock-box to one minute, and leaves the room. The timer countdown is displayed on a screen built into the lid. Participants are encouraged to spend this waiting time by rolling the die “to practice.”

Participants are instructed to roll the die *exactly once* as soon as the box unlocks, and write down the outcome on a now-accessible slip of paper. Clearly, since there is no one else in the room, participants have an opportunity to lie, i.e., misreport the outcome of the final die roll. Participants are then asked to place the die back in the pouch before leaving the room, so as to leave no trace of their actual outcome. Outside, they proceed to the payoff station and submit their outcome slip to the experimenter. Having written down their outcome in private, participants need not lie to the experimenter in conversation.

We implement the one-minute wait for two reasons. First, some participants may be suspicious about the fairness of the die. This waiting time allows them to check whether the die is fair. Second, and more importantly, a time delay serves as a mandatory period of deliberation, which could increase participants' awareness of the opportunity to lie (Lohse et al., 2018) or override their intuition to cooperate with the experimenter, i.e., be honest (Rand et al., 2014). In other words, the delay helps ensure participants think rationally about their incentives and response. Indeed, in line with Fischbacher and Föllmi-Heusi (2013), we make every effort to convey that dishonesty could not possibly be observed and therefore could never be punished, *without* ever explicitly saying so in order to avoid priming.

### 2.3 Conducting the experiment

Our experiment was conducted over 17 days in March 2019 at four sites within the National University of Singapore. A total of 524 participants (60% female) took part. We solicited demographic information from participants, summarized in Appendix B.2. We required at least 94 participants per treatment group to detect a 0.5 difference in mean reported dice roll outcome from the expected outcome of 3.5, with 80% power and 5% significance level (see Appendix A.6 for calculation). All our treatment groups had at least 101 participants.

The experiment was advertised via online research recruitment websites, mass emails from university admin personnel, and posters at high-traffic hubs around campus. Participants could sign up online or simply walk in to the experiment site. On average, the experiment took 10 minutes per participant, who were paid a show up fee of \$5 SGD ( $\sim$  4USD). We precluded participants from participating more than once.<sup>5</sup>

Upon arrival, participants randomly drew a unique ID, which also determined their treat-

---

<sup>5</sup>As we collected no personally identifying information, we could not systematically check names or IDs. As precautions, the same experimenters conducted all sessions, and participants were all informed of the preclusion condition before they were assigned their groups.

ment group. Once assigned to a group, participants (except the charity-only group) were informed that they could earn additional payoff based on their outcomes during the experiment. Since the additional payoffs varied by treatment group, the exact amount of potential bonus earnings was not specified. Participants were not given information on other treatment conditions, and were prohibited from communicating verbally or with their cellphones.

Participants received written instructions (Appendix A.1), a payoff chart showing how their reported outcome corresponded to additional payoffs (Appendix A.3), and, with the exception of the self-only group, a menu of charities. Participants' questions were raised and addressed privately. They were then individually brought inside a private room to carry out the die-rolling task described in Section 2.2. Upon completion, participants proceeded to a private payoff station, where an experimenter paid participants (in cash) and/or charities (via instant online bank transfer) accordingly. Participants were requested to oversee the calculation and disbursement of payoffs, and to verify all bank transfers. After leaving the payoff station, participants privately answered a questionnaire (Appendix A.4), which was later matched to their reported outcomes. Finally, subjects were debriefed and paid their show-up fee.

### 3 Results

Table 2 below shows the shares of participants in our self-only group who reported each outcome, compared with those in Shen et al. (2016), a die-roll experiment also run in Singapore, and the meta-study by Abeler et al. (2019). Results are similar.<sup>6</sup> Both Fischbacher and Föllmi-Heusi (2013) and Shen et al. (2016) have participants do other unrelated tasks before and after the die-roll task, respectively. Our results show that excluding these “extra” experiments has no effect on participants' behavior.

#### 3.1 Maximal lying and pure honesty occur in all groups.

Figure 1 depicts graphically the share of participants in each treatment group who reported each possible payoff, the numerical frequencies for which are reported in Table 3.

As 6 is the highest value on a die, we call misreporting a 6 *maximal lying*. If all participants were honest, we would expect one-sixth of them to report having rolled a 6. However, in every treatment group, 6 is reported significantly more frequently than expected — even in the

---

<sup>6</sup>We regress reported outcome against various demographic characteristics and find no results of significance. See Appendix B.4 for details.

| Group                                 | Share of participants (%) |      |      |      |      |      |        |
|---------------------------------------|---------------------------|------|------|------|------|------|--------|
|                                       | 1                         | 2    | 3    | 4    | 5    | 6    | $N$    |
| This study<br>(self-only group)       | 4.6                       | 12.0 | 15.7 | 17.6 | 17.6 | 32.4 | 108    |
| Singapore study<br>Shen et al. (2016) | 9.8                       | 8.3  | 11.2 | 14.1 | 27.8 | 28.8 | 205    |
| Meta-study<br>Abeler et al. (2019)    | 7.6                       | 8.4  | 10.9 | 15.6 | 25.0 | 32.5 | 30,185 |

Table 2: Shares of participants in self-only group who reported each payoff, compared to a similar study in the same country (Shen et al., 2016), and a meta-study of 90 experiments (Abeler et al., 2019)

charity-only group, where participants cannot increase their own payoff (one-tailed binomial tests,  $p < 0.05$ ). The frequency of reported 6s decreases monotonically as the share to charity increases. While reported 6s are twice as frequent as expected for the self-only and 90-self groups, this ratio decreases to 1.5 in charity-only.

If, following Fischbacher and Föllmi-Heusi (2013), we assume that participants do not lie to decrease their payoff, then participants who report a 1 must be *purely honest*. As seen in Table 3, a nonzero share of participants in every treatment group do report 1s. This share is smallest in the self-only group (4.63%), largest in charity-only (17.65%), and consistent across split-payoff groups (11-12%).

| Group        | Share of participants (%) |         |        |       |          |          |     |
|--------------|---------------------------|---------|--------|-------|----------|----------|-----|
|              | 1                         | 2       | 3      | 4     | 5        | 6        | $N$ |
| Self-only    | 4.63***                   | 12.04   | 15.74  | 17.59 | 17.59    | 32.41+++ | 108 |
| 90-self      | 11.88                     | 11.88   | 9.90** | 18.81 | 15.84    | 31.68+++ | 101 |
| 50-self      | 11.82                     | 10.00** | 11.82  | 16.36 | 21.82+   | 28.18+++ | 110 |
| 10-self      | 10.68*                    | 7.77*** | 10.68* | 18.45 | 27.18+++ | 25.24++  | 103 |
| Charity-only | 17.65                     | 11.76   | 11.76  | 18.63 | 15.69    | 24.51++  | 102 |

\* (+)  $p < 0.1$ , \*\* (++)  $p < 0.05$ , \*\*\* (+++)  $p < 0.01$ .

Table 3: Share of participants in each treatment group who reported each payoff. Stars (plus signs) indicate significance for one-sided binomial tests that observed share is smaller (larger) than expected share of 16.67%

### 3.2 Partial lying increases as take-home share decreases.

In line with Fischbacher and Föllmi-Heusi (2013), we define *partial lying* as a lie where participants misreport to increase payoff, but do not claim the maximum possible payoff (e.g., report a

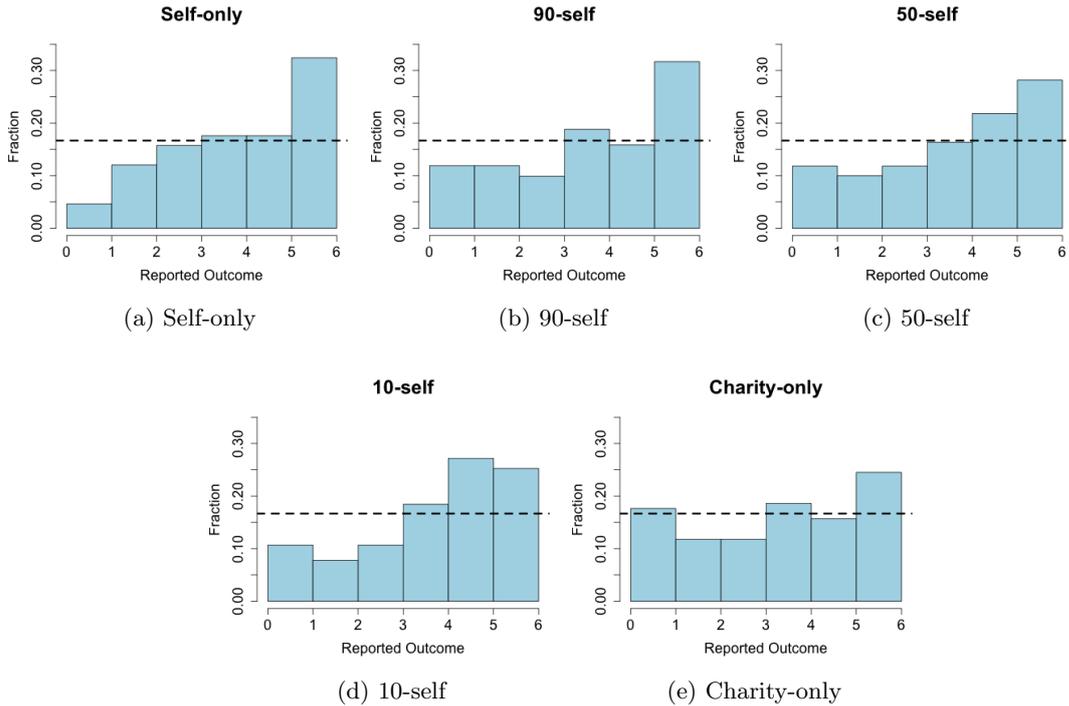


Figure 1: Share of participants in each treatment group who reported each payoff

5 instead of a 6). Table 3 shows that as take-home share decreases from 100% to 10%, the share of reported 6s decreases while the share of reported 5s increases. This suggests that participants switch from maximal to partial lying as their stake in the payoff decreases. Partial lying seems to disappear in the charity-only group, where the share of reported 5s drops to the expected share.

### 3.3 Overall lying behavior is significant in all groups except charity-only.

If all participants were honest, we would expect the mean outcome from repeatedly rolling a die to be 3.5. We find that the mean reported outcome significantly differs from 3.5 for all groups *except* charity-only (two-sided, one-sample t-tests,  $p < 0.001$ ). In the charity-only group we find no such significance ( $p = 0.14$ ).

We formally test the reported outcome distributions of each treatment group against discrete uniform distributions of comparable sample size. We conduct two non-parametric tests: the Kolmogorov-Smirnov (K-S) test, which examines the single largest vertical difference between two distributions, and the Wilcoxon rank-sum (WRS) test, which pools the two distributions and tests for clustering. Because our data is (i) discrete and (ii) has multiple ties across samples to be compared, the standard versions of both tests are known to produce conservative  $p$ -values.

We use Arnold and Emerson (2011)’s K-S Monte Carlo  $p$ -value simulation method using 1,000 replicates, and Marx et al. (2016)’s dynamic programming WRS solution, and report both accurate and conservative  $p$ -values in Table 4.

| Group        | Kolmogorov-Smirnov test<br>(two-sample, one-sided) |              | Wilcoxon rank-sum test<br>(two-sided) |              |
|--------------|--|--------------|---------------------------------------|--------------|
|              | Accurate   | Conservative | Accurate                              | Conservative |
| Self-only    | 0.000***   | 0.002***     | 0.001***                              | 0.001***     |
| 90-self      | 0.002***   | 0.009***     | 0.012**                               | 0.013**      |
| 50-self      | 0.000***   | 0.004***     | 0.008***                              | 0.008***     |
| 10-self      | 0.000***   | 0.000***     | 0.003***                              | 0.003***     |
| Charity-only | 0.196  | 0.405        | 0.270                                 | 0.270        |

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 4: Nonparametric tests of observed outcome distributions against uniform. Accurate and conservative  $p$ -values reported for comparison

We observe similar results under both tests. Reported outcome distributions deviate significantly from the uniform in self-only and all split-payoff groups. Once participants do not benefit from the lie, however, overall lying behavior diminishes drastically, to the point that we cannot reject the null that the charity-only group’s reported outcome distribution is uniform. Chi-squared goodness-of-fit tests and Fisher exact tests yield similar results (Appendix B.1).

Furthermore, we follow Abeler et al. (2019) in estimating the share of participants who lied in each treatment group. Assuming that participants who actually did roll a 5 or 6 have no incentive to further inflate their results, we expect two-thirds of participants to have had incentive to lie, i.e., those who rolled 1, 2, 3, or 4. For each treatment group, we compute the difference between the reported and expected frequencies of  $\{1,2,3,4\}$ , and then divide by the expected frequency. This yields the estimated share of participants in each group who lied, conditional on having had incentive to do so. As seen in the first column of Table 5, all treatment groups have remarkably similar estimated lying shares of around 25% — once again with the exception of charity-only, with an estimated lying share of 10%.

### 3.4 Self-only participants are most likely to lie about having lied.

We also asked participants through a post-task questionnaire (Appendix A.4) whether they were honest in their report. 90.4% of all participants said “Yes”, 4.3% said “No”, and 5.2% did not respond. The share of each treatment group who admitted to lying — i.e., responded “Yes” — is reported in the second column of Table 5. Participants answered this questionnaire before

| Group        | Share of participants (%) |                |                        |
|--------------|---------------------------|----------------|------------------------|
|              | Estimated lying           | Admitted lying | Admitted/Estimated (%) |
| Self-only    | 25.00                     | 0.93           | 3.72                   |
| 90-self      | 21.30                     | 4.95           | 23.24                  |
| 50-self      | 25.00                     | 5.45           | 21.80                  |
| 10-self      | 28.63                     | 7.77           | 27.14                  |
| Charity-only | 10.30                     | 2.94           | 28.54                  |

Table 5: Discrepancies between estimated and self-admitted shares of lying. Estimated lying is computed as  $1 - \frac{\text{reported \# of } \{1,2,3,4\}}{\text{expected \# of } \{1,2,3,4\}}$ . That is, if, as expected, 2/3 of respondents report  $\{1,2,3,4\}$ , this value is 0, and, if none report  $\{1,2,3,4\}$ , this value is 1

receiving their show up fee, but after receiving their game payoff. While dishonest participants might have reason to conceal their lie, it is hard to imagine honest participants falsely claiming to having told a lie. Hence, self-admitted lying can be considered a lower bound on the number of misreports.

As seen in Table 5, the estimated lying share is greater than self-admitted share across treatment groups — evidence that participants lie about lying. We report the percentage of estimated liars who admitted to having lied in the third column of Table 5. We find that only 3.72% of estimated liars in the self-only group admit to having done so, compared to 22-29% in all other treatment groups. Indeed, sharing *any portion* of the payoff with charity seems to make it easier for participants to admit having lied.

Perhaps participants who did not answer the question on lying are all liars, and their very non-response is an admission of guilt. Our finding is robust to considering non-respondents as self-admitted liars (Appendix B.3).

## 4 Discussion

As shown in Table 2, our results for the self-only group are comparable to previous findings in the literature. We are therefore confident that our experimental design is not inherently biased — or at least, is biased in the same way as most previous studies.

Previous research has primarily focused on the two “extreme” groups, self-only and other-only (e.g., charity-only). In our self-only group, we find that maximal lying and pure honesty occur at frequencies comparable with Fischbacher and Föllmi-Heusi (2013) and Gneezy et al. (2018), whose experimental designs closely resemble ours. In our charity-only group, we find the

lowest occurrence of maximal lying and the highest occurrence of pure honesty. This result is consistent with Wiltermuth (2011) and Klein et al. (2017), who find significantly less lying when it only benefits other participants,<sup>7</sup> but not with Lupoli et al. (2017), who find that participants lie for charity approximately as much as for themselves.<sup>8</sup>

Notably, our experiment allows us to study how lying behavior changes as we gradually move along these two extremes. We find that estimated lying share (Table 5) remains virtually unchanged as charity share increases: reducing private material gains from 100% to 90%, 50%, or even 10% has no apparent effect. However, reducing private material gains from 10% to 0% results in a 15 percentage point reduction in lying rates. Indeed, if we look instead at overall lying behavior as captured by reported outcome distributions, we find that lying is significant in all groups except charity-only (Table 4). Overall, participants seem insensitive to the *size* of their private payoff (intensive margin), but sensitive to its *existence* (extensive margin). One plausible explanation is that participants believe charities would not want to receive “dirty money” that had been obtained fraudulently. This belief might be bypassed, however, when participants get to take home even as little as 10%, by thinking that in this case, “everyone would agree” that it is acceptable to lie to get more for oneself.

Our results are consistent with Klein et al. (2017), who also implement a range of split-payoff treatments. Even though their “other” is not a charity but an anonymous participant, we find a similar pattern of results: lying frequency holds steady when participants receive all or most of the payoff, but sharply decreases when they receive little or none of it.<sup>9,10</sup> Our results, however, differ from Wiltermuth (2011) and Gino et al. (2013), who find significantly greater lying rates when payoffs are split between self and an anonymous other participant than when payoffs are solely for self or other.

What could explain the inconsistent results about lying when payoffs are split? One possible reason is that the identity of the recipient may indeed matter. When payoffs are split and the

---

<sup>7</sup>We note that Klein et al. (2017) use a within subject design, i.e., all subjects were asked for their choices in each of the treatments.

<sup>8</sup>Maggian (2019) also has a die-rolling experiment about lying for charity, but it is not directly comparable with ours. While her control group is similar to our self-only group (reporting an “odd” number yields €4, reporting “even” yields €1), participants in her charity group must face a tradeoff: all money they claim is taken away from a donation to a charity (while for the control group, all money participants claim is taken away from the researcher). That is, reporting “even” would give €1 to the participant and the remaining €3 to charity. Hence, in order to increase donations, participants must *under-report*. She finds no difference between the two groups.

<sup>9</sup>They estimate the following lying rates in each (\$ for self, \$ for other) group: (\$5,\$0): 25%; (\$4,\$2): 35%; (\$3,\$4): 26%; (\$2,\$6): 26%; (\$1,\$8): 15%; (\$0,\$10): 15%.

<sup>10</sup>Erat and Gneezy (2012) also find that propensity to lie diminishes from 65% to 49% as the portion shared with an anonymous other participant diminishes from 50% to 9%.

beneficiary is another individual, lying typically increases. However, Maggian (2019) suggests that people are more willing to act unethically against an organization — even a charitable one — than against an individual. This would imply that lying to benefit an organization might provide less moral flexibility than lying to benefit an individual. This might explain why Wiltermuth (2011) and Gino et al. (2013) find that the rate of lying peaks when fellow research participants also benefit, but we observe no such peak when splits are with charity.

Another possible reason is that the amount of effort implicit in the lying task may matter. Both Wiltermuth (2011) and Gino et al. (2013) ask participants to privately perform a semi-skilled task (e.g., solve matrices, unscramble anagrams) and report their performance, whereas Klein et al. (2017) and our study assign participants virtually effortless tasks (toss a coin, roll a die). This suggests that when reported scores implicitly signal about participants’ effort or ability, participants might inflate scores when payoffs are split to impress the “other” participant or the experimenter.

Furthermore, in a single-blind study like ours, where participants report their results individually, participants might believe that experimenters update their beliefs about participants’ honesty, and behave accordingly. As Abeler et al. (2019) show, two types of lying costs are consistent with stylized results in the literature: a preference for *being* moral (self-image), and a preference for *being seen* as moral (social image/reputational cost). The latter usually serves to encourage socially applauded behavior when it can be directly observed or indirectly inferred, as in our study. Indeed, it has been shown that in dictator games, participants tend to donate more in single-blind than double-blind studies (Franzen & Pointner, 2012; Hoffman et al., 1996).<sup>11</sup> This suggests that reputational lying costs might be *smaller* for prosocial lying, as compared to selfish lying. In other words, being perceived as a liar might carry a smaller social punishment if it were done for charitable reasons.

However, we find instead that lying decreases sharply when payoffs go entirely to charity. This does not mean the conjecture is untrue: for instance, we find that *lying about lying* sharply increases when payoffs go entirely to the self (Table 5). This suggests the reputational costs matter: *admitting* to lying is less costly when some share of payoff goes to charity, regardless of the size of that share. Thus, we take our findings to suggest that, on average, the decrease in reputational lying costs is simply insufficient to warrant lying more for charity than for oneself.

---

<sup>11</sup>In die-rolling honesty games, Fischbacher and Föllmi-Heusi (2013) find no significant difference in lying between single- and double-blind studies. That said, their study has no charitable-donation component, i.e., all payoffs are for self only.

To the best of our knowledge, no study comparing single- and double-blind studies has yet been done for prosocial lying. Future research in this direction should shed more light on the issue.<sup>12</sup>

There are three final considerations that could influence our findings: (i) participant disinterest in the available charities; (ii) lack of trust that experimenters would actually make the donations; and (iii) charities acting as an *ex ante* nudge. We discuss them briefly in turn.

First, perhaps participants did not really care about the charity the money goes to. In that case, those in the charity-only group would have no incentive to lie, as their utility gain from lying would be essentially non-existent. However, we do observe maximal lying in that particular group. While it is possible that a handful of respondents did not care about any of the charities, we believe that the majority did, as we populated our menu of charities with a diverse cross-section of social issues with wide appeal (Appendix A.5).

Second, perhaps participants did not trust experimenters to actually make the donations. We believe this is not the case: as soon as participants received task instructions, they were made aware that the donation process would happen online, in their presence, before they could leave the experiment site (Appendix A.1). The menu of charities provided to participants contained the official websites, logos and mission statements of each charity, to demonstrate their legitimacy. Our experiment was also supported and approved by the National University of Singapore, which is considered a reputable institution in the country. Furthermore, donations were indeed made in front of participants as promised, so word-of-mouth could not have hurt the reputation of our experiment. For all these reasons, we believe that our results could not have been affected by participants' mistrust of experimenters.

Third, perhaps charitable donation served as a nudge for honest behavior: since most participants were forced to think about charities before performing the die-rolling task, their willingness to lie might have been affected *ex ante*. If this were true, then we would expect participants in the self-only group, who were unaware of the purpose of the experiment, to lie significantly more than all other participants. We find that this is not the case: participants in all split-payoff groups exhibit significant lying behavior (Result 3.3) and lie at rates comparable to the self-only group (Table 5). However, as stated above, the "nudge" might work when the *only* beneficiary of the lie is the charity itself: when payoffs are split, then selfish considerations overcome the

---

<sup>12</sup>A double-blind version of our die-rolling task would have a group of participants enter the experiment room together, roll the die in randomly-chosen private cubicles, collect money from envelopes left at each cubicle, and leave the room together. Experimenters would thus be unable to link specific individuals to specific envelopes. For more details on double-blind implementations, see Hoffman et al. (1996) and Fischbacher and Föllmi-Heusi (2013).

“nudge.”

## 5 Conclusion

This study aims to assess how undetectable lying behavior changes depending on the beneficiary of the lie: oneself, a charity of choice, or both oneself *and* a charity. We can summarize our findings thus. First, there exist participants in each group who lie to inflate their reported number. Second, participants lie more often when their share of the payoff is strictly positive. Third, participants admit to lying more often when the charity’s share of the payoff is strictly positive. Our findings are mostly consistent with previous studies, with differences likely arising from the identity of the recipient and the nature of the task. Future work could further investigate the effect of reputational lying costs by comparing single- and double-blind versions of our prosocial lying paradigm. As ours is the first study on prosocial lying not carried out in a “Western” country, we cannot disregard the possibility that differences in our results are due to prevailing cultural norms in Southeast Asia. Future work could investigate the extent to which such norms mediate prosocial lying behavior.

## References

- Abeler, J., Nosenzo, D., & Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, *87*(4), 1115–1153.
- Andreoni, J. (2006). Philanthropy (S.-C. Kolm & J. M. Ythier, Eds.). In S.-C. Kolm & J. M. Ythier (Eds.), *Handbook of the Economics of Giving, Altruism and Reciprocity (Volume 2)*. Elsevier B.V.
- Arnold, T., & Emerson, J. (2011). Nonparametric goodness-of-fit tests for discrete null distributions. *The R Journal*, *3*(2), 34–39.
- Engel, C. (2011). Dictator games: A meta study. *Experimental Economics*, *14*(4), 583–610.
- Erat, S., & Gneezy, U. (2012). White lies. *Management Science*, *58*(4), 723–733.
- Fischbacher, U., & Föllmi-Heusi, F. (2013). Lies in disguise: An experimental study on cheating. *Journal of the European Economic Association*, *11*(3), 525–547.
- Franzen, A., & Pointner, S. (2012). Anonymity in the dictator game revisited. *Journal of Economic Behavior & Organization*, *81*(1), 74–81.
- Gino, F., Ayal, S., & Ariely, D. (2013). Self-serving altruism? The lure of unethical actions that benefit others. *Journal of Economic Behavior & Organization*, *93*, 285–292.
- Gneezy, U., Kajackaite, A., & Sobel, J. (2018). Lying aversion and the size of the lie. *American Economic Review*, *108*(2), 419–453.
- Hermann, D., & Mußhoff, O. (2019). I might be a liar, but I am not a thief: An experimental distinction between the moral costs of lying and stealing. *Journal of Economic Behavior & Organization*, *163*, 135–139.
- Hoffman, E., McCabe, K., & Smith, V. L. (1996). Social distance and other-regarding behavior in dictator games. *The American economic review*, *86*(3), 653–660.
- Klein, S. A., Thielmann, I., Hilbig, B. E., & Zettler, I. (2017). Between me and we: The importance of self-profit versus social justifiability for ethical decision making. *Judgment and Decision Making*, *12*(6), 563–571.
- Konow, J. (2010). Mixed feelings: Theories of and evidence on giving. *Journal of Public Economics*, *94*(3-4), 279–297.
- Levitt, S. D., & List, J. A. (2007). What do laboratory experiments measuring social preferences reveal about the real world? *Journal of Economic Perspectives*, *21*(2), 153–174.

- Lohse, T., Simon, S., & Konrad, K. (2018). Deception under time pressure: Conscious decision or a problem of awareness? *Journal of Economic Behavior & Organization*, *146*, 31–42.
- Lupoli, M. J., Jampol, L., & Oveis, C. (2017). Lying because we care: Compassion increases prosocial lying. *Journal of Experimental Psychology*, *146*(7), 1026–1042.
- Maggian, V. (2019). Negative externalities of cheating: An experiment with charities (A. Bucciol & N. Montinari, Eds.). In A. Bucciol & N. Montinari (Eds.), *Dishonesty in Behavioral Economics*. Academic Press.
- Marx, A., Backes, C., Meese, E., Lenhof, H.-p., & Keller, A. (2016). EDISON-WMW: Exact dynamic programming solution of the Wilcoxon–Mann–Whitney Test. *Genomics, Proteomics & Bioinformatics*, *14*(1), 55–61.
- 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).
- Shen, Q., Teo, M., Winter, E., Hart, E., Chew, S. H., & Ebstein, R. P. (2016). To cheat or not to cheat: Tryptophan hydroxylase 2 snp variants contribute to dishonest behavior. *Frontiers in behavioral neuroscience*, *10*, 82.
- Wiltermuth, S. (2011). Cheating more when the spoils are split. *Organizational Behavior and Human Decision Processes*, *115*, 157–168.