Matteo Ploner

Lie for me: An experiment about delegation, efficiency, and morality

CEEL Working Paper 2-22

Cognitive and Experimental Economics Laboratory

Via Inama, 5 38100 Trento, Italy

http://www-ceel.economia.unitn.it
tel. +39.461.282313

Lie for me:
An experiment about delegation, efficiency, and morality

Matteo Ploner*

September 15, 2022

Abstract

Individuals and organizations may delegate others to perform actions they would not do themselves because of moral constraints. In our experimental setting, a Principal can either self-report a value in a die-under-the-cup task or delegate the report to an Agent who has no material interest in the report. We experimentally manipulate the relative efficiency of the report: the Agent's prospect either stochastically dominates that of the Principal or vice versa. We find that Principals have a high propensity to lie and delegate only when the Agent's prospect is more efficient. Agents generally behave honestly, but those with higher prosocial attitudes tend to lie when assigned an inefficient lottery, most likely not to let down the Principal.

Keywords: Honesty, Decision Making for Others, Belief-based preferences.

JEL Codes: D91, C91, D81

* *Corresponding author*
University of Trento, Department of Economics and Management
Via Inama, 5
38123 Trento (Italy)
matteo.ploner@unitn.it

1 Introduction

Though they may deliver desirable outcomes, there are actions that we refrain from undertaking because of moral considerations. In these circumstances, actions are not guided by strict consequentialist logic but by general concerns about the fairness of the process leading to an outcome. For example, an individual concerned about animal welfare might refrain from eating animal-derived food if they had to farm it. However, the same individual may enjoy scrambled eggs served at the campus cafeteria when lunching with colleagues. Other examples may involve unsustainable waste disposal, non-cruelty-free productions, and exploitative (child) labor. Here, we investigate how delegation may alleviate moral constraints and provide a "moral wiggle-room" to the decision-maker.

In our experimental setting, individuals observe a binary signal, whose realization is associated with a high or low payoff, and freely decide which realization to report and earn the associated payoff. The signal is private and undetected by any external observer (die-under-the-cup paradigm, see Fischbacher and Föllmi-Heusi, 2013). Thus, mere payoff-maximization considerations would always dictate reporting the realization associated with the high payoff. However, moral considerations may encourage reporting truthfully. The decision-maker (Principal) has the possibility of not choosing firsthand and delegating another individual (Agent) to choose for them, to avoid the moral dilemma. When delegation happens, the Agent, who has no direct stake in the decision, can freely choose the report's content. A false report is generally more advantageous for the Principal than a truthful report. The Agent's setting is similar to the one faced by the Principal, with a strictly private signal, not observable to any external observer, Principal included. In such a setting, Agents may decide to lie to favor their Principal altruistically.

Our focus is on the Agents' choices when delegated to choose for their Principals. To study under which circumstances altruistic lies may emerge, we experimentally manipulate the relative payoffs associated with the high/low signal: the Agent's prospect is either more or less efficient than the Principal's. Furthermore, we also allow Agents to ask for a fee for the service provided to Principals. Finally, we control for the social orientation of Agents. The intuition supporting our controlled variations is that Agents are more likely to lie in favor of the Principal when delegation implies firmer Principals' expectations of deceitful behavior by the Agent. Our data partially support this intuition by showing that Agents with stronger prosocial attitudes tend to be dishonest only when delegation is disadvantageous for the Principal. At the same time, Principals correctly anticipate that Agents' dishonesty is not generalized and decide not to delegate and, untruthfully, report firsthand.

1.1 Literature Review

Experimental evidence shows that people are ready to lie, at least partially, to achieve higher economic outcomes (see the review in Abeler et al., 2019). An important stylized fact that emerges from this now vast literature is that pecuniary incentives matter for lying (Gneezy, 2005) but are not the sole determinant of moral behavior, with several contextual and individual factors affecting the decision to lie (Gerlach et al., 2019). While several theoretical models have been put forward to explain this pattern, one that has attracted a lot of attention incorporates economic incentives to lie, intrinsic costs of lying—*wanting to be honest*—and image concerns—*wanting to appear honest* (see, Khalmetski and Sliwka, 2019; Abeler et al., 2019; Gneezy et al., 2018).

In our experiment, participants can delegate the choice to lie to another participant. We interpret this as a way to save on the psychological cost of lying. Individuals may use delegation as a "moral wiggle room" (Dana et al., 2007) to appear honest to others and themselves. One could argue that asking others to perform a potentially unethical action is already a sign of dishonesty. However, individuals seem to systematically manipulate their perception of the world to fit their motivations (Epley and Gilovich, 2016). This kind of motivated reasoning translates into self-deception, taking the form of strategic ignorance, reality denial, and self-signaling (Bénabou and Tirole, 2016). In this perspective, "outsourcing" choices in ethical dilemmas may result from motivated reasoning to reduce the psychological costs associated with lying. Empirical support to the ”diffusion of responsibility” when delegating others to perform an immoral act comes from Hamman et al. (2010) and Bartling and Fischbacher (2011).

Together with the motives of those who delegate, we study here the choices of those who are delegated. When delegated, the decision to behave unethically has direct consequences only for the party that delegates. Lies of the surrogate liars have thus a notable prosocial aspect. Levine and Schweitzer (2014) portray the moral character as made of two main components: honesty and benevolence. The authors argue that the latter may prevail when the two conflict. Indeed, individuals who lie to benefit others in ethical dilemmas are perceived to be more moral than individuals acting honestly. Erat and Gneezy (2012) also investigate the tension between honesty and prosociality and find that about 1/3 of the participants in the experiment are ready to lie to help others, even if they have to pay a (small) cost. At the same time, some participants refrain from lying when this would be advantageous for all parties involved. Weisel and Shalvi (2015) investigate cooperation in dishonest behavior and show that not only do couples coordinate on dishonest actions when this is beneficial for both but also when it is beneficial only for one of the two. In a lab-in-the-field experiment, Michailidou and Rotondi (2019) find that the lower the social distance between people, the higher the propensity to lie prosocially. Buckle et al. (2021) compare lying for others and lying for self in a die-under-the-cup experiment and find that though a considerable share of participants lies for others without any direct monetary incentive to do so, lying for others is less likely than lying for self.

According to Levine and Schweitzer (2015), prosocial lies likely originate in the desire to help others in a consequential logic aimed at improving others’ payoffs. However, as already noted by Gneezy (2005), prosocial lies may also originate in nonconsequential logic that considers the process leading to a certain outcome and not only the outcome itself. Belief-based preferences (Battigalli and Dufwenberg, 2022) and, specifically, guilt-aversion (Battigalli and Dufwenberg, 2007) may provide a useful theoretical support to understand prosocial lying. The intuition is that a subject asked to lie to benefit another subject will lie because they believe that the other expects them to lie. When this is the case, the decision to lie prosocially is not driven by payoff considerations but by the second-order beliefs of the surrogate liar. Dufwenberg and Dufwenberg (2018) present a theoretical investigation of belief-based preferences in moral dilemmas. Still, the role of this family of preferences in surrogate lying is widely under-investigated.

The experimental work closer to ours is Erat (2013), implementing a sender-receiver game (Erat and Gneezy, 2012) with the possibility for a sender to "hire" an agent to send the message to a receiver. The incentives of the sender and the agent are perfectly aligned, and both are better off when they mislead the receiver than when the receiver correctly guesses the signal. At the same time, being misled is costly for the receiver. The author finds that about one-third of senders decide to hire an agent to send the message, and hiring is more likely when the damage to the receiver is more considerable. Kandul and Kirchkamp (2018) builds upon Erat and Gneezy by investigating the truth-telling preferences of senders who delegate and testing whether delegation choices affect behavior in a subsequent independent donation task. The authors confirm that many participants delegate others to lie and show that many of those who delegate would have told the truth if they had to send the signal themselves. This evidence supports the conjecture that many perceive delegation as a way to earn more without paying the psychological cost of lying.

The present work innovates on the literature on delegation in unethical tasks in several dimensions. First, the agents have no material incentive to lie for their principal. Second, outcomes of the unethical task may differ when the principal firsthand makes a choice or when the choice is delegated to an agent. As illustrated below, this manipulation is crucial to test our hypothesis that the principal's expectations affect the agent's choices. Third, agents can ask for a fee to become surrogate liars in one of our conditions. This option makes the delegation closer to a market transaction. Fourth, we control for social attitudes of agents, a potential mediator for dishonest behavior (Soraperra et al., 2019).

2 Method

2.1 Tasks and Roles

Main Task

Participants are randomly matched in couples and assigned to two roles: *Principal* and *Agent*. The Principal in the couple can decide whether to retain control and participate in a task or give up control and delegate the Agent to participate. The general structure of the task is the same for the Principal and the Agent: a six-faced fair die is privately tossed, and the self-reported outcome of the die defines the earnings in a lottery (*die-under-the-cup paradigm*, see Fischbacher and Föllmi-Heusi, 2013). Precisely, an even outcome delivers a high return in the lottery, while an odd outcome delivers a low one. The key feature of the task is that the task participant is not bound to report the actual outcome and can decide to lie, knowing that the experimenter will not know the result of the die toss.

The Principal is endowed with a *Lottery A* and the Agent with a *Lottery B*. The main difference between the Principal and the Agent's lotteries is the size of the rewards associated with alternative outcomes. As illustrated by Table 1, the High outcome (H) of the Principal's Lottery A is always €7, and the Low outcome (L) is always €3. Thus, when the Principal truthfully reports the outcome, the lottery has an expected value of €5. Differently, the Agent's faces three distinct Lotteries B that change in the size of H and L: in Lottery $B\downarrow$, H=€2 and L=€6; in Lottery $B\rightarrow$, H= €3 and L= €7; in Lottery $B\uparrow$, H= €4 and L= €9. Importantly, given these payoffs and assuming fair reports, the Agent's Lotteries B can thus be ranked in terms of stochastic dominance: $B\uparrow \succ B\rightarrow \succ B\downarrow$.

	Principal	Agent		
Outcome	A	$B\downarrow$	$B\rightarrow$	$B\uparrow$
High	7	6	7	9
Low	3	2	3	4

Table 1: Lottery outcomes (€)

The three lotteries are administered to Agents over three distinct rounds. In each round, an Agent is randomly matched with a Principal and reports the die's outcome for one of the three lotteries. The display order of the Agent's lotteries is random within a matching group to control for possible order effects. As specified above, the Principal is endowed with Lottery A in all three rounds. Thus, the following combinations of lotteries are presented to each participant in random order: $A|B \rightarrow$, $A|B \uparrow$, $A|B \downarrow$.

In each round, the interaction between the Principal and the Agent follows this protocol:

- the Principal and the Agent are both informed about the Lottery A's and B's nature (common knowledge);
- the Principal decides whether to play Lottery A (retain control) or ask the Agent to play Lottery B (delegate);
  - the Principal rolls the die: if the Principal chooses to retain control (Lottery A), the self-reported outcome is relevant for the definition of their earnings; if the Principal chooses to delegate (Lottery B), the rolling of the die and the self-report are not relevant for the definition of payoffs.[1]
- before knowing the choice of the Principal, the Agent privately rolls the die and self-reports the outcome for the lottery (strategy method).[2]

For the Principal, the earnings in the task are given by the outcome in the lottery, either A, when retaining control or B, when delegating. The actual payment is then given by one of the three rounds randomly chosen. For the Agent, the actual payment is given by a fixed amount that is not affected by choices in the experiment. For both types of participants, actual earnings are announced only at the end of the experiment.

[1]This procedure was implemented to keep Principals busy and avoid potential inference by the Agents about Principals' actual choices, as the rolling of the die is noisy.

[2]We adopt a strategy method to collect choices of the Agents to increase the power of our data, as it is reasonable to anticipate that actual delegation would differ across lottery configurations.

Ancillary Tasks

In each round, after the lottery task, participants are asked about their beliefs about the actions of their counterparts. Specifically, Principals are asked about the likelihood that Agents will report an even outcome, and Agents are asked about the likelihood that the Principal will choose Lottery A. Participants can choose among five probabilities (0%, 25%, 50%, 75%, and 100%) associated to the choice of the counterpart (see Appendix A for the elicitation interface). For each probability, participants see the earnings when the two options available to the counterpart are actually chosen. This way, incentives to truthfully reveal are directly presented to participants without the need to illustrate to them the underlying scoring rule.[3]

At the end of the three rounds of the main task, participants are asked to participate in a series of Social Value Orientation (SVO) allocation tasks (Murphy et al., 2011). In each task, they are asked to choose the preferred allocation of money between themselves and another participant, randomly chosen among other participants in the room. Different allocation tasks present different combinations of payoffs for the decision-maker and the recipient. The joint analysis of choices allows us to classify participants into alternative social types.[4] After collecting all choices, half of the participants are randomly assigned to the role of decision-maker and the other half to the role of recipient. Couples are made, and payoffs are computed according to one of the randomly chosen choices of the decision-maker.

[3] We use the following quadratic scoring rule to elicit probabilities about our binary outcome: $S_i(p) = \alpha - \beta \sum_{k=1}^{2}(I_k - p_k)^2$, where $\alpha, \beta > 0$ and $I_k = 1$, when event k is realized and equal to 0 when even k is not realized (e.g. Palfrey and Wang, 2009). In our implementation, $\alpha = 2$ and $\beta = 1/10000$, delivering earnings between €0 and €2 according to the probability associated with the event realized.

[4] Participants took part in a series of independent allocation tasks; in each task, they had to choose an allocation to them and an anonymous other among nine alternatives. The combination of payoffs for the two parties changed in every task. After everyone had chosen, half of the participants would randomly become the actual allocator. One of their choices would define the payment to them and the matched recipient selected from among those who are not allocators. Social types are then computed following the procedure described by Murphy et al. (2011).

2.2 Treatments

The main controlled manipulation in our paper is the combination of lotteries A and B. The three alternative combinations, $A|B$, $A|B\downarrow$, $A|B\uparrow$ allow us to study delegation dynamics under different economic incentives (see Section 3 for a description of our theoretical framework). In addition to this within-subjects manipulation, we also introduce a between-subjects manipulation. In the *NoFee* condition, the interaction follows the protocol described above, with Principals choosing to delegate or retain control and Agents self-reporting the outcome of the die's roll. Differently, in the *Fee* condition, the Agents can ask to be paid a fee to roll the die (F). Specifically, the Agent communicates the requested fee to the Principal, and then the Principal decides whether to delegate or retain control. If the Principal delegates, they will pay the fee; otherwise, they will not. Then, the Agent will self-report the outcome of the die, knowing that if the Principal delegates, they will earn the fee in addition to the fixed amount. Notably, the fee is bounded ($0 \leq F \leq 1$) to preserve the stochastic dominance ordering of Agent's lotteries.

In addition to the sessions involving only human participants, we ran a series of control sessions matching human Principals with a computer algorithm programmed to play the lotteries fairly (CPU). The general structure of the interaction is the same as in the *NoFee* condition, with the only difference being that the die roll is performed virtually by the computer acting in the role of Agent. The Principal is made fully aware of the behavior of the algorithm. These sessions aim to obtain a benchmark for the Principal's behavior when the counterpart is bound to behave honestly.

2.3 Participants and Procedures

The experiment was conducted at the Cognitive and Experimental Economics Laboratory (CEEL) of the University of Trento (Italy). The computerized part of the experiment was administered and programmed in z-Tree (Fischbacher, 2007). Each participant was randomly assigned a cubicle with a desktop computer, and participants' choices were collected and

Treatment	Role	N
CPU	Principal	46
Fee	Agent	45
	Principal	45
NoFee	Agent	45
	Principal	45
Total		226

Table 2: Participants

processed by the software. In each cubicle, the participants found a die and a cup to roll the die privately, without being observed by the experimenter and other participants. Once all were seated, participants were given a few minutes to read instructions privately. Afterward, a member of the staff read instructions aloud (see Appendix A for a translated version of the instructions). Participants could ask questions and clarifications by raising their hands. Answers were generally given in private.

A total of 226 individuals took part in the experiment (see Table 2 for a summary), over 12 sessions: 90 in the *Fee* condition, 90 in the *NoFee* condition, and 46 in the *CPU* condition. In the *Fee* and *NoFee* conditions, half of the participants were randomly assigned to the role of Principal and half to the role of Agent. Participants were associated in matching groups of size 6, and the order of lotteries B was randomized at the matching group level. In the *CPU* conditions, human participants were matched with a computer algorithm.

The experiment lasted about 90 minutes, including the initial procedures and payments in cash. On top of the €3 show-up fee, Principals earned €8.67 and Agents €7.81, on average. The earnings of the Principals are given by the outcome of the lottery, deducted by the fee paid to the Agent, and by the outcome in the belief elicitation task. The earnings of the Agents are given by the fixed payment of €7, the fee obtained from the Principal, and the outcome in the belief elicitation task.

3 Theoretical Framework

To derive our research hypotheses, we present here a theoretical framework that leverages on recent contributions about the role of emotions in social interactions (Battigalli and Dufwenberg, 2022). Specifically, we focus on the psychological cost of lying originating in self-image concerns and on guilt experienced when letting-down the other.

In the self-report task, the Principal privately observes the true signal $\omega \in \{L, H\}$, with $H > L$ and $p(L) = p(H) = .5$. However, the reported value $r \in \{L, H\}$ can be freely chosen by the Principal, irrespective of the actual ω. We assume that when the report is selfishly misaligned relative to the true signal, Principals experience a cost for lying that originates in self-image concerns. The utility function of a Principal j, borrowed from Abeler et al. (2019), is thus given by

$$u_j(r_j, c(r_j, \omega), \theta) = \pi(r_j) - \theta_j c(r, \omega)$$

where $\pi(r_j)$ is the payoff associated with the reported value by P, with $\pi(H) = H$ and $\pi(L) = L$; $c(r, \omega)$ captures the psychological cost of lying, with $c(H, L) > 0$ and $c = 0$ in all other configurations of ω and r; $\theta_j \geq 0$ is the weight given to the cost of lying: the larger θ the larger the psychological cost of lying in terms of self-image.[5] The expected utility of a given lottery is equal to $E[u(r_j) = p_L u_j(\omega = L) + p_H u_j(\omega = H)$. Given that the cost of lying is possibly experienced only for $\omega = L$, $u_j(\omega = H) = H$, it follows that the Principal j will report $r = H|\omega = L$ (misreport) when $\theta < \frac{H-L}{c(H,L)}$. For lower values of θ, misreporting becomes more likely, keeping fixed the cost of lying and the lottery spread. In the extreme case of $\theta_j = 0$, the Principal will always report $r = H$, even when $\omega = L$.

An important assumption about Principal's utility is that the cost of lying is experienced only when the lie is due to actions under direct control of the decision-maker, namely a

[5] As convincingly shown by Abeler et al. (2019), this simple model can be extended to encompass psychological costs associated with reputation towards an audience. However, for our presentation, we focus on self-image concerns only.

misreport. However, when the lie is attributable to another individual, no psychological cost is experience. This assumptions is justified in light of the literature on moral-wiggle rooms showing that simple manipulations in the connection between actions and consequences can suffice to remove moral concerns (e.g., Bénabou and Tirole, 2016). Within this framework, the act of delegating is a way for the Principal to preserve their self-image and avoid the costs of lying associated to misbehavior. We denote $\alpha_j^{r|\omega}$ the beliefs of j about the report r by Agent i conditional on $\omega \in \{L, H\}$. The expected utility of delegating i to choose is given by $E[u(r_i)] = p_L\big(\alpha_j^{L|L}L + (1-\alpha_j^{L|L})H\big) + p_H\big(\alpha_j^{L|H}L + (1-\alpha_j^{L|H})H\big)$. We assume that the Agent will possibly lie only to favor the principals and, thus, $\alpha_j^{L|H} = 0$. A Principal will thus deliberately choose to delegate when $E[u(r_i)] > u_j(r_j, c(r_j, \omega), \theta)$.

We now turn to the Agents' utility function, which is made of a constant monetary component, of the psychological cost of lying and, differently than the Principal, of a psychological component related to simple guilt aversion (Battigalli and Dufwenberg, 2007). This captures the intuition that an Agent delegated to choose "incorporates" the Principal's beliefs into their utility function and aims to behave in line with these anticipated beliefs not to disappoint the Principal. Thus, the utility function of an Agent i is given by

$$u_i(c(r,\omega), \theta, \gamma, \alpha_p) = \pi_i - \theta_i c(r,\omega) - \gamma_i \big[E[\pi_j; \alpha_j] - \pi_j(r)\big]^+$$

where $E[\pi_j; \alpha_j]$ is the payoff Principal j expects from the interaction, given their first-order beliefs α_j and $\pi_j(r)$ is the payoff j actually gets, given i's report r; the parameter $\gamma \geq 0$ captures i's sensitivity to let-down j, i.e. to deliver a payoff that is lower than what the Principal expects.

Given that the payoff of an Agent is not a function of their report, they must choose their report r to minimize the impact of the psychological cost of lying, taking into account both self-image concerns and guilt aversion, to minimize $\theta_i c(r,\omega) + \gamma_i \big[E[\pi_j; \alpha_j] - \pi_j(r)\big]^+$. It is essential to notice that for anticipated optimistic beliefs of the Principal, there is a trade-off between the two components: to meet the Principal's expectations, one must lie in half of the cases, and this implies a cost in terms of self-image. Within this framework, the choice of an Agent is a maximization process that depends on own sensitivity to image concerns, guilt aversion, and anticipated beliefs of the counterpart.

3.1 Research Hypotheses

We derive here a few research hypotheses, within the theoretical framework sketched above. We start with our predictions for Agents. First, it is important to consider that an Agent who shows no guilt sensitivity ($\gamma = 0$) has no incentive to misreport the outcome of the die roll. Furthermore, guilt sensitivity induces an Agent to misreport an unfavorable outcome for the Principal only when it overcomes the cost of lying, as captured by θ. Here, we do not attempt to estimate γ and θ parameters, but we derive some research hypotheses about the impact of alternative experimental conditions on the likelihood of misreporting the outcome of the die roll. Our first set of hypotheses refers to Agents (Hypotheses 1, 2, and 3). The controlled manipulations in the experiment aim at altering the relevance of belief-based preferences for Agents, keeping constant the psychological cost of lying. Hypothesis 4 focuses on Principals and on the possibility of "saving" on psychological costs of lying by delegating the choice to the Agent.

We assume that when a Principal delegates an Agent for Lottery $B \downarrow$, the expectations of an advantageous (mis)report are manifest to the Agent. The lottery is indeed stochastically dominated by A, and the only reasonable justification for delegation is the assumption that the Agent will advantageously misreport. For other lotteries the intentions behind delegation are less manifest as, in expected terms, $B \rightarrow$ and $B \uparrow$ are equal or better than A, respectively. Accordingly, we assume that it is easier for the Agent to intercept the Principal's expectations in Lottery $B \downarrow$. This will make let-down aversion more salient for $B \downarrow$ than for other B lotteries, even though first-order beliefs of the Principal are never directly communicated to the Agent.

Hypothesis 1 (Agents: Prospects) *Agents are more likely to advantageously misreport in lottery $B \downarrow$ than in other B lotteries.*

Our second hypothesis also originates in belief-based preferences and refers to the possibility for the Agent to ask for a fee to roll the die. Here, we interpret the request of a fee by the Agent as an implicit commitment signal likely to raise the Principal's expectations ($\alpha_j^{H|L}$), similarly to the non-binding promise of reciprocation in Battigalli and Dufwenberg (2007). Accordingly, an Agent who is concerned about not letting down the Principal should be more likely to advantageously misreport when they ask for a fee than when not.

Hypothesis 2 (Agents: Fee) *Agents are more likely to advantageously misreport when they can ask a fee (Fee) than when they cannot ($NoFee$).*

In our belief-based framework, a crucial role is played by the parameter γ, capturing sensitivity to let-down aversion. This parameter is highly idiosyncratic, but we speculate here that the prosocial attitudes of individuals are correlated with γ. Expressly, we assume that prosocial individuals are generally more concerned about others' welfare than selfish individuals and, thus, are more inclined to experience guilt.

Hypothesis 3 (Agents: Prosociality) *Prosocial agents are more likely to misreport than selfish agents*

Delegation may represent an opportunity for Principals to avoid experiencing the moral costs of lying (θ) and still enjoy the material benefits of lotteries. In our framework, the decision to delegate requires the comparison between the expected pecuniary outcomes of the lottery and the psychological cost of lying. Under the assumption of honesty by Agents, the difference between the expected outcome for $B \uparrow$, $B \rightarrow$ and $B \downarrow$, and the best outcome of Lottery A are €0.5, €-1.5 and €-2.5, respectively. Given these monetary incentives, we predict that delegation will be more frequent for $B \uparrow$ than for other lotteries. At the same time, Principals may anticipate that Agents' behavior is affected by experimental conditions, as detailed by our Hypothesis 1. Thus, we predict that delegation will be more frequent for $B\downarrow$, when guilt-aversion is more likely to play a role, than for Lottery $B\rightarrow$. Finally, for the mechanism detailed in Hypothesis 2, we also predict that delegation is more likely when a fee is paid than when it is not.

Hypothesis 4 (Principals: Delegation) *Principals will delegate, especially for Lotteries $B\uparrow$ and $B\downarrow$ and when they pay a fee.*

4 Results

We first analyze Agents' behavior across alternative experimental conditions (Section 4.1) and then turn to delegation choices by Agents (Section 4.2). The main results are first presented with the support of graphical illustrations and statistical tests and then corroborated by regression analysis.

4.1 Agents

Of the 270 Agents' choices we collected across the experimental conditions, 161 (59.5%) reported an even outcome.[6] A non-parametric test shows that the individual-level overall frequency of an even report is not significantly different from the fair benchmark report of .5 (p-value=0.296).[7] Overall, Agents in our experiment seem to faithfully report the outcome of the private toss of the die.

[6]Moshagen and Hilbig (2017) suggest that the share of dishonest reports is likely to be underestimated when focusing on actual outcomes because a share of honest participants will also observe the favorable outcome. They suggest correcting this and estimating the share of dishonesty as $d = \frac{q-p}{1-p}$, where q is the empirically observed frequency of a favorable outcome, and p is the baseline probability of observing it (in our case, $d = 0.193$). However, this estimate, and the associated statistics, hold when dishonest participants only lie favorably. In our context, we prefer to abstain from this strong assumption and focus on empirically observed frequencies, even if this may underestimate the share of dishonest participants.

[7]We take the individual-level overall frequency across prospects—how many times they report an even outcome out of the three choices—as the unit of analysis for our tests to control for repeated choices across prospect types. For paired comparison, we adopt a two-tailed Wilcoxon Signed Rank test (WSR), while for unpaired comparisons, we adopt a two-tailed Wilcoxon Rank Sum test (WRS).

Lottery Type

Figure 1 presents the self-reported outcomes for the three B lotteries separately.

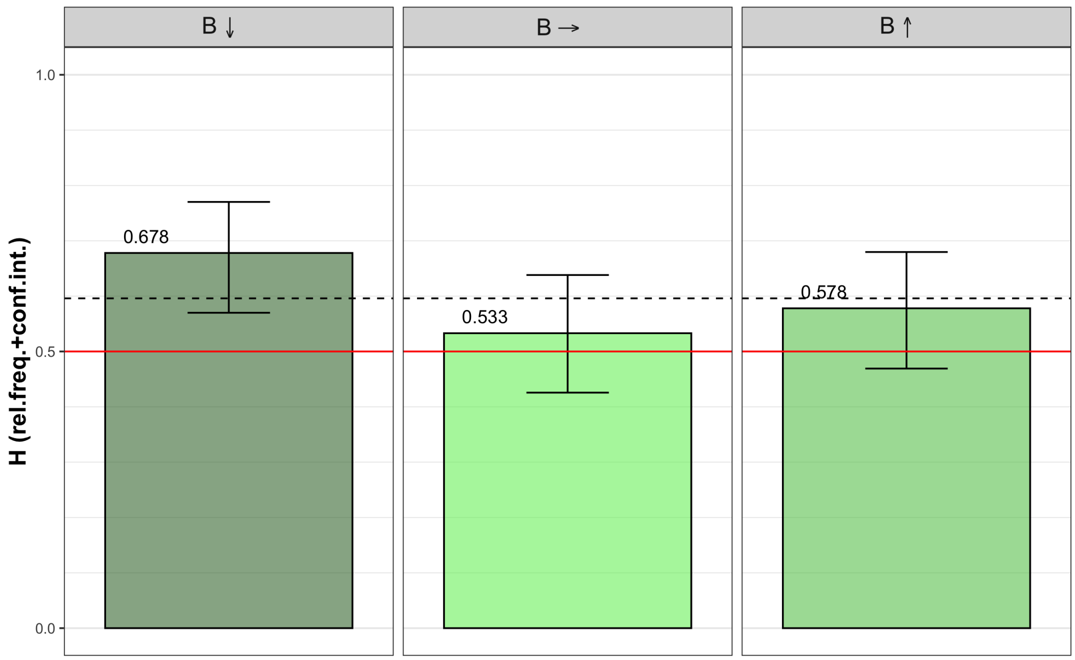

Figure 1: Agents' report by prospect type

As Figure 1 shows, the frequency of even reports is above the .50 benchmark in all prospects. However, as shown by the confidence interval of a proportion test, only in prospect $B\downarrow$, can we reject the null hypothesis that the frequency of even reports is equal to the fair benchmark outcome.

Table 3 summarizes the output of a generalized linear mixed model (GLMM, Logit link) controlling for repeated observations at the individual level (for each of the 90 agents we collect 3 choices). The dependent variable is given by Agent's report (Even=1, Odd =0). As explanatory variables, we adopt dummy variables capturing the type of lottery, with lottery $B\rightarrow$ as the baseline.

Table 3 highlights the statistically significant difference between $B\downarrow$ and the baseline lottery, with a higher likelihood of observing a H outcome in the former. A linear hypothesis test on outcome shows that there is not significant difference between $B\downarrow$ and $B\uparrow$ (χ^2 test,

$H \sim$	
(Intercept)	0.156 (0.246)
$B\downarrow$	0.712 (0.337)*
$B\uparrow$	0.212 (0.325)

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\circ}p < 0.1$

Table 3: Lottery Types and Agents' report (GLMM, Logit link)

p-value=0.136). Results from the regression analysis and the descriptive analysis reported above lead us to formulate our first result, which provides support to Hypothesis 1

Result 1 *Advantageous reports by Agents are significantly more frequent than those expected under honesty assumptions in Lottery $B\downarrow$. Furthermore, for this lottery, advantageous outcomes are significantly more frequent than in the baseline Lottery B.*

Fees

We now turn to the Agents' behavior across distinct fee conditions. Table 4 provides us with some descriptive statistics about fee requested by Agents in the *Fee* condition.

Lottery type	Mean	SD	Median
$B\rightarrow$	0.624	0.342	0.7
$B\downarrow$	0.638	0.317	0.6
$B\uparrow$	0.651	0.289	0.6

Table 4: Fees requested (Summary statistics)

Descriptive statistics are quite similar across different lottery types, as confirmed by non-parametric tests that do not identify any significant difference in fees requested across lottery types (WSR, all p-values$>$ 0.587). To investigate whether the fee requested is a function of the intention to report an H outcome, we also compare fee levels among those reporting different outcomes. Differences between those reporting H and those reporting L are very small and equal to 0.009, 0.120, and 0.008 in lottery $B\rightarrow$, $B\downarrow$, and $B\uparrow$, respectively.

Non-parametric tests confirm that these differences are not significantly different than zero (WRS, all p-values> 0.220). Thus, fees seem not to be affected by the type of lottery or by the subsequent report of the Agent.[8]

Concerning reports across Fee conditions, the Agents tend to report an H outcome more frequently under the *NoFee* than under the *Fee* condition, 65.9%, and 53.3%, respectively. A Wilcoxon rank-sum test (WRS) shows that the individual-level overall frequency of an even report is marginally different between the two conditions (p-value=0.080). Figure 2 provides a detailed representation of the frequency of H outcomes for different prospects and fee conditions.

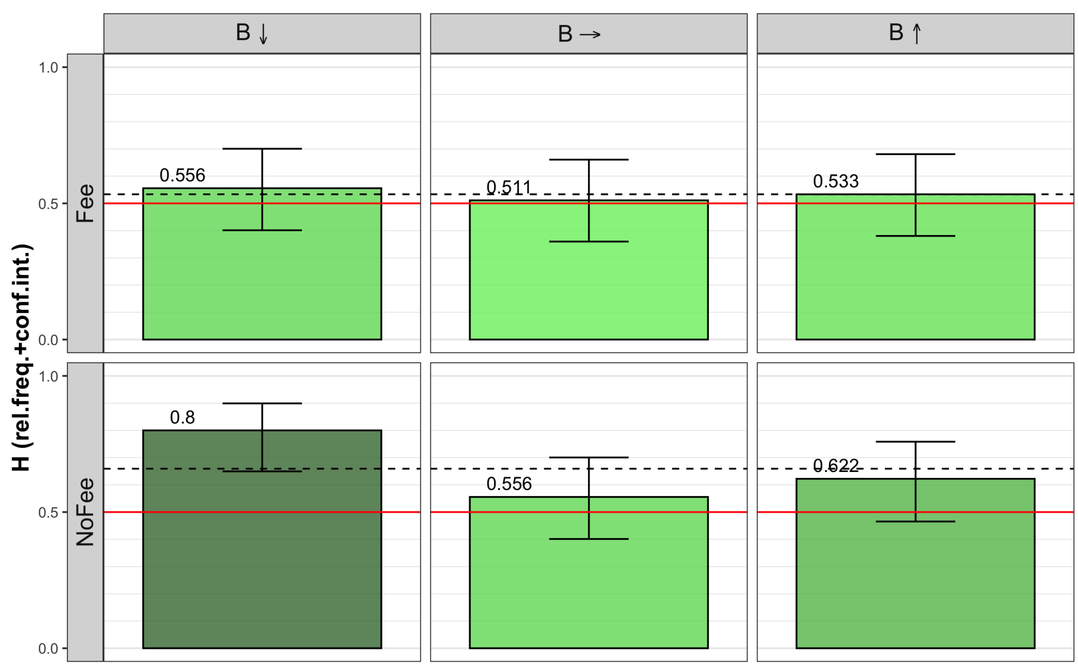

Figure 2: Agents' report by prospect type and fee treatment

The H frequency is larger for every prospect in the *NoFee* than in the *Fee* condition. Notably, in condition *NoFee* for prospect $B\downarrow$, the frequency is significantly different from the fair benchmark, according to the proportion test whose confidence intervals are reported

[8] That the levels of fee requested do not affect the reported outcome is also confirmed by a regression analysis (GLMM, Link logit) showing that the estimated coefficient of the fee level is not statistically significant (coeff.=0.321; std.err.= 0.634).

in the graph.

The GLMM output of Table 5 shows that the likelihood of an H outcome is significantly higher in lottery $B\downarrow$ than in the baseline $B\rightarrow$, even controlling for the fee treatment (model 1). Furthermore, H reports are marginally more likely in the *NoFee* condition. In model 2, we also consider the interaction between the lottery type and the Fee treatment. As the table shows, no significant interaction between the two treatment variables is observed; however, a linear hypothesis test shows that the likelihood of an even outcome is significantly higher in the $NoFee$ condition than in the Fee condition for the $B\downarrow$ lottery (χ^2 test, p-value=0.010) but not for the $B\uparrow$ lottery (χ^2 test, p-value=0.487).

$H \sim$	(1)	(2)
(Intercept)	-0.151 (0.291)	0.051 (0.345)
$B\downarrow$	0.713 (0.337)*	0.209 (0.457)
$B\uparrow$	0.212 (0.325)	0.104 (0.456)
NoFee	0.616 (0.323)$^{\circ}$	0.209 (0.490)
$B\downarrow \times$ NoFee		1.120 (0.687)
$B\uparrow \times$ NoFee		0.218 (0.650)

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\circ}p < 0.1$

Table 5: Fees and Agents' report (GLMM, Logit link)

The evidence collected about the impact of fees in reports leads us to formulate the following result, which goes against our hypothesis 2

Result 2 *Advantageous reports by Agents are marginally more likely in the* NoFee *condition than in the* Fee *condition, especially for lottery $B\downarrow$.*

Social Types

As detailed above, the independent SVO task allows us to classify our participants into different social types. Specifically, 48 Agents out of 90 (53.3%) are classified as *Individualist* and 42 (46.7%) as *Prosocial*. The relative frequency of H reports is 61.8% for the former and 57.1% for the latter. A WRS shows that the individual-level overall frequency of a H

report is not significantly different between the two conditions (p-value=0.375). Figure 3 provides a detailed representation of the frequency of H reports for different prospects and social types.

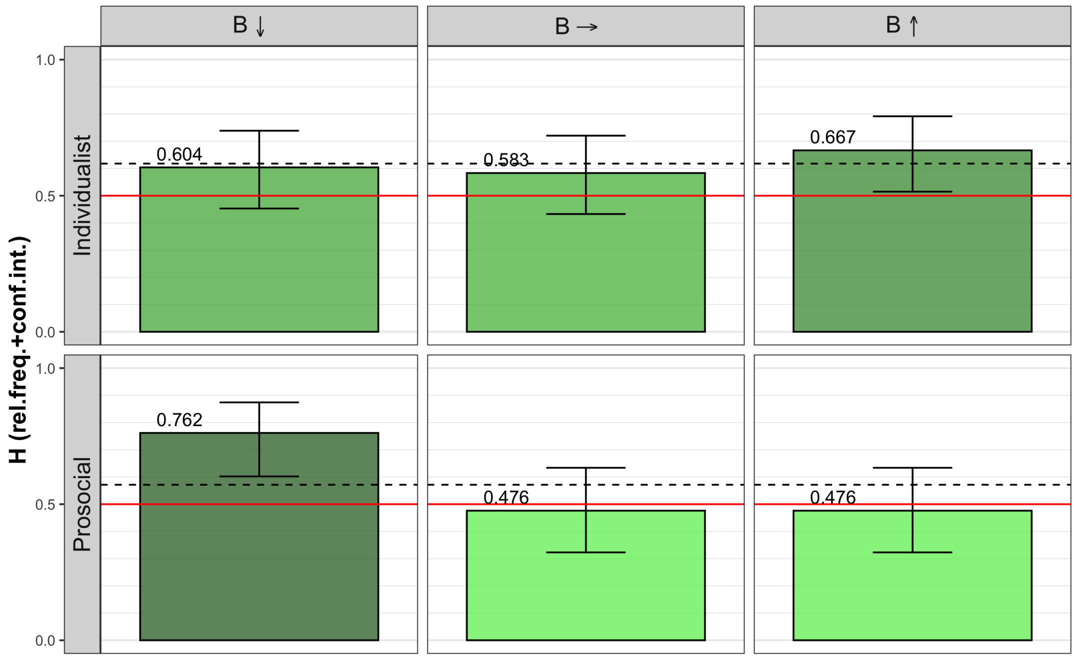

Figure 3: Agents' report by prospect type and social value orientation

The graph highlights different patterns of behavior for the two social types: the *Prosocial* over-report the number of H events for prospect $B\downarrow$ (one-sample proportion test, p-value=0.001); the Individualist over-report for Prospect $B\uparrow$ (one-sample proportion test, p-value=0.030).

The GLMM output of Table 6 highlights that in $B\downarrow$ there is a significantly larger likelihood of reporting H than in the baseline $B\rightarrow$ lottery, controlling for social type (column 1). Importantly, the social type classification does not have a significant impact on the likelihood of a H report. Column 2 completes evidence from column 1 by showing that the impact of $B\downarrow$ is mainly due to the interaction between prosocial attitudes and lottery $B\downarrow$. In line with evidence from Figure 3, a linear hypothesis test shows that the likelihood of a H outcome in $B\uparrow$ is smaller for the Prosocial than for the Individualist, though only marginally so (χ^2 test, p-value=0.068).

$H \sim$	(1)	(2)
(Intercept)	0.266 (0.290)	0.404 (0.348)
$B \downarrow$	0.712 (0.337)*	0.104 (0.454)
$B \uparrow$	0.212 (0.325)	0.427 (0.463)
Prosocial	−0.234 (0.324)	−0.517 (0.505)
$B \downarrow \times$ Prosocial		1.369 (0.692)*
$B \uparrow \times$ Prosocial		−0.427 (0.664)

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\circ}p < 0.1$

Table 6: Social types and Agents' report (GLMM, Logit link)

Result 3 *Overall, different social types do not differ in the likelihood of reporting an advantageous outcome. However, Prosocial agents display a higher propensity to misreport for lottery $B \downarrow$. In contrast, Individualist agents tend to misreport more in $B \uparrow$.*

4.2 Principals

Delegation

Principals largely retain control over the lotteries and delegate the Agent to choose only in 38% of the cases. Figure 4 provides a refined picture of delegation choices across experimental conditions and lottery types.

The figure shows that the delegation frequency differs significantly from the randomization benchmark identified by the continuous red line in the figure (one-sample proportion tests, all p-values < 0.026). However, opposite patterns emerge for lotteries $B \downarrow$ and $B \rightarrow$, where Principals do not generally delegate, and $B \uparrow$, where they mostly delegate. This suggests that the expected value of the lottery is the main driver of he Principal's decision to delegate.

To understand the Principals' motivations when choosing to delegate, we analyze the (incentivized) beliefs about the expected choices of the Agents. Figure 5 illustrates the distribution of beliefs about an Even report (H) across alternative experimental conditions

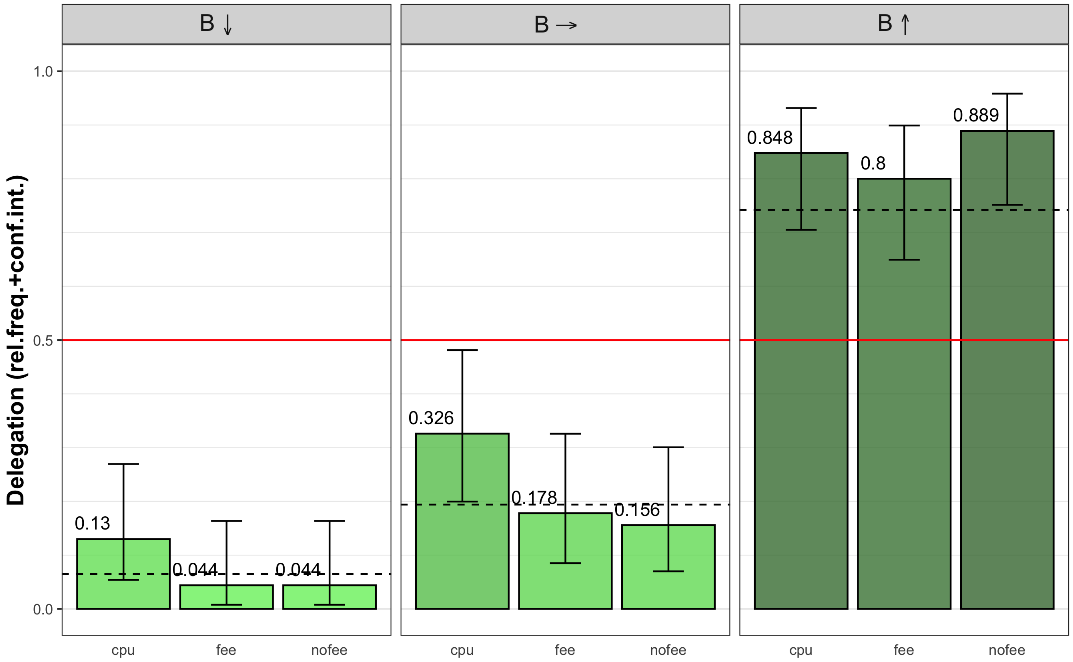

Figure 4: Principal's decision to delegate

and lotteries. The vertical dashed line captures the average belief of the Principal, and the continuous red line captures the observed average frequency by Agents in that condition: the closer the lines, the more accurate the average beliefs of the Principals.

The figure clearly shows that principals correctly anticipate the generalized truthful reporting of Agents, with the distribution mode in correspondence to 50%. The only exception is in condition *NoFee* for Lottery $B\downarrow$, where Principals do not anticipate the significant deviation from honest reporting by Agents. Correctly, the Principals do not see delegation as a profitable choice, given the anticipated honesty of Agents.

Table 7 reports the output of a GLMM estimation, which takes as a dependent variable the decision to delegate (1=delegate, 0=retain control). As explanatory variables, we consider lottery types (the baseline is $B\rightarrow$), fee treatment (the baseline is CPU), beliefs about the likelihood of an even report by the Agent, and Prosocial type. In column (2), we report a model estimating the potential interaction between fee treatment and lottery type.

The regression outcomes in column (1) highlight a higher likelihood of delegating for

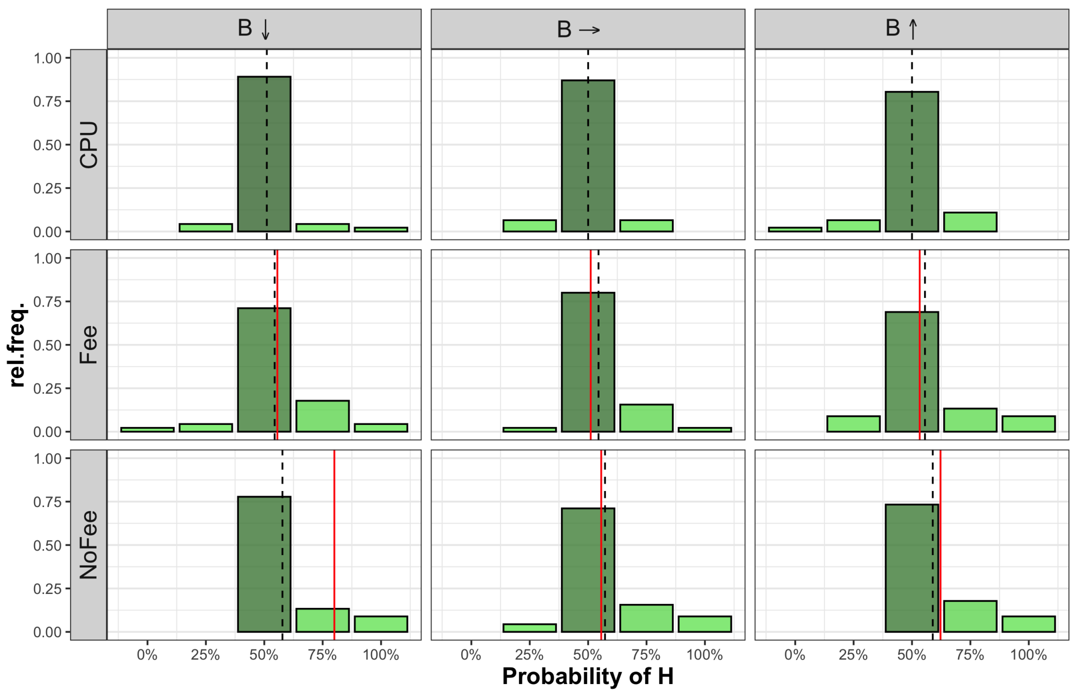

Figure 5: Principal's decision to delegate

	(1)	(2)
(Intercept)	−1.334 (1.077)	−1.067 (1.121)
$B\downarrow$	−1.616 (0.470)***	−1.602 (0.666)*
$B\uparrow$	4.341 (0.711)***	3.587 (0.814)***
Fee	−0.949 (0.523)°	−1.154 (0.720)
NoFee	−0.685 (0.518)	−1.338 (0.736)°
Belief	−0.112 (0.319)	−0.133 (0.329)
Prosocial	0.816 (0.428)°	0.832 (0.443)°
$B\downarrow \times$ Fee		−0.313 (1.116)
$B\uparrow \times$ Fee		0.817 (0.934)
$B\downarrow \times$ NoFee		0.016 (1.090)
$B\uparrow \times$ NoFee		1.840 (1.017)°

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\circ}p < 0.1$

Table 7: Principals' choice to delegate (GLMM, Logit link)

$B\uparrow$ and a lower likelihood for $B\downarrow$, relative to the baseline $B\rightarrow$. Furthermore, there is a marginally lower likelihood of delegating in the Fee condition, and Prosocial types are marginally more likely to delegate. Outcomes in column (2) do not highlight any significant interaction between the lottery types and the fee treatment condition.

To complete our description of Principals' motivations, we provide a description of Principals who do not delegate their choice to the Agent. Of the 253 Principals who retain control over the report, 201 (79.4%) report a H outcome. A non-parametric test on individual level overall frequency shows that the frequency of H outcomes is significantly higher than the fair benchmark of .5 (WSR, p-value=0.006). The frequency of H reports is equal to 0.817, 0.736, and 0.952 for lottery $B\downarrow$, $B\rightarrow$, and $B\uparrow$, respectively. One-sample proportions tests confirm that these frequencies all are significantly different from the fair benchmark of .5 (all p-values < 0.001). Thus, Principals display a stronger propensity to misreport than Agents.

Result 4 *Principals are not reluctant to advantageously misreport and, anticipating the honest behavior of Agents, delegate more for lottery* $B\uparrow$ *and less for lottery* $B\downarrow$ *than for lottery* $B\rightarrow$.

Finally, to get a measure of the profitability of delegation, we compute the average payoff for each condition and lottery type, when the choice is made by the Agent. The difference between these averages and the actual payoff of those retaining control gives us a measure of the profitability of delegation. Table 8 reports the results of such analysis. As the Table shows, delegation always generates a loss, on average, but in condition $NoFee$ for Lottery $B\uparrow$.

Treatment	Lottery type	Average Profit
Fee	$B\rightarrow$	-1.391
Fee	$B\downarrow$	-2.857
Fee	$B\uparrow$	-0.984
NoFee	$B\rightarrow$	-1.041
NoFee	$B\downarrow$	-1.149
NoFee	$B\uparrow$	0.111

Table 8: Average Profits of Delegation

5 Discussion and Conclusions

Previous works have shown that delegation may constitute a strategy to "shift the blame". Our results partly differ from this conclusion by showing that the Principals prefer to retain control over the decision-making process and misreport the private signal observed to their advantage. This opportunistic strategy turns out to be profitable, given the overall honesty of the Agents. Thus, Principals seem not to generally perceive delegation as an expedient to avoid moral concerns, like shown in previous studies (Erat, 2013). It might be that the distance between intentions and consequences created by delegation is not enough, and Principals prefer to "pay" the psychological cost of deception rather than incur the opportunity cost and the risk embedded in the delegation to an honest agent. When looking closely into our data, we can notice that in lottery $B \uparrow$ delegation is persistent, even though Principals correctly anticipate the general honesty of agents. In this lottery, a dishonest report by the Principal warrants Eur 7 for sure, while the expected value of honest reports by Agents is Eur 6.5. Thus, delegation here has a monetary cost and also entails risk. From this, we could infer that delegation may serve to avoid moral costs, but these costs seem quite low for our Principals, as shown by the lack of delegation in the other lotteries.

On the Agent's side, we predicted that more deceitful reports should be observed for lottery $B \downarrow$ than for other lotteries, given the implicit Principal's expectations of a dishonest report embedded in the delegation of $B \downarrow$. Indeed, our results show that $B \downarrow$ is the only lottery in which there is a significant deviation from the fair benchmark. Moreover, the frequency of favorable outcomes in this lottery is significantly higher than that of the baseline lottery $B \rightarrow$. At a closer investigation, it emerges that this effect is mainly driven by the behavior of prosocial Agents, in line with our hypothesis 3. Those who are classified as prosocial are more likely to react to the nature of the prospect in line with our hypotheses based on guilt aversion. Thus, we argue that prosocial types are more reactive to the implicit message sent by the Principal when delegating the dominated lottery. The connection between prosociality and sensitivity to guilt has not yet been systematically investigated and future works should devote attention to this nexus.

Overall, Agents' behavior seems compatible with the belief-based explanation outlined by our hypotheses. Alternative explanations appear less prominent in light of collected data. If Agents were simply selfish, no specific pattern should have been observed in the data, as selfish players have no direct stake in their decisions. If Agents were altruistic, they should have over-reported the favorable outcome in every prospect. If they were experiencing a cost for lying, they should have reported honestly for every prospect, irrespective of the underlying payoff structure of the Principal's and Agent's lottery. Misreports may result from a trade-off between altruistic attitudes and lying costs. In this perspective, the monetary incentives for lying are lower for $B \downarrow$, as the highest payoff is way lower than that for $B \uparrow$. Thus, we would have expected more lying in $B \uparrow$ than in $B \downarrow$, given the psychological cost of lying. Curiously, this is what we observe in $B \uparrow$ for the individualist types, although the effect is moderate. Thus, while the prosocial seem to react to indirectly revealed expectations of the counterpart, individualists react more to the monetary consequences of lying.

While differences across prospects provide overall support to the hypotheses based on guilt aversion, the fee manipulation is in contrast to our hypothesis. We speculated that the request for a fee would have provided a signal of commitment, similar to the promise in Charness and Dufwenberg (2006), triggering a more robust mechanism of let-down aversion. However, we find precisely the opposite: the possibility of requesting a fee reduces misreports. In hindsight, this can be possibly due to a crowding out of the intrinsic motivations to lie by Agents (Ryan and Deci, 2000). Once a fee is paid, altruistic lying loses its moral value and becomes just a violation of a deontological norm. Given that the fee obtained in return is quite small, it could be that the monetary compensation did not overcome the loss in this intrinsic dimension.

Many real-life situations allow us to delegate others to perform actions that we would deem immoral. Among others are unsustainable waste disposal, non-cruelty-free productions, and exploitative (child) labor. The experimental evidence reported here shows that delegation mechanisms may potentially trigger cooperative deception and the adoption of inferior technology only when the "vicious" expectations of the delegating party are apparent and the delegated party is motivated by prosocial concerns. At the same time, delegation is undertaken only when its expected monetary cost is low. Thus, the scope of delegation as a conscious-clearing mechanism seems quite limited and likely to foster cooperative deception only under specific circumstances.

References

Abeler, J., Nosenzo, D., and Raymond, C. (2019). Preferences for Truth-Telling. *Econometrica*, 87(4):1115–1153. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA14673.

Bartling, B. and Fischbacher, U. (2011). Shifting the blame: On delegation and responsibility. *The Review of Economic Studies*, 79(1):67–87. Publisher: Oxford University Press.

Battigalli, P. and Dufwenberg, M. (2007). Guilt in Games. *The American Economic Review*, 97(2):170–176. Publisher: American Economic Association.

Battigalli, P. and Dufwenberg, M. (2022). Belief-Dependent Motivations and Psychological Game Theory. *Journal of Economic Literature*, 60(3):833–882.

Buckle, G. E., Füllbrunn, S., and Luhan, W. J. (2021). Lying for others: The impact of agency on misreporting. *Economics Letters*, 198:109677.

Bénabou, R. and Tirole, J. (2016). Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives*, 30(3):141–164.

Charness, G. and Dufwenberg, M. (2006). Promises and Partnership. *Econometrica*, 74(6):1579–1601. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0262.2006.00719.x.

Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1):67–80.

Dufwenberg, M. and Dufwenberg, M. A. (2018). Lies in disguise: A theoretical analysis of cheating. *Journal of Economic Theory*, 175:248–264.

Epley, N. and Gilovich, T. (2016). The Mechanics of Motivated Reasoning. *The Journal of Economic Perspectives*, 30(3):133–140. Publisher: American Economic Association.

Erat, S. (2013). Avoiding lying: The case of delegated deception. *Journal of Economic Behavior & Organization*, 93:273–278. Publisher: Elsevier.

Erat, S. and Gneezy, U. (2012). White Lies. *Management Science*, 58(4):723–733.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Ex-*

perimental Economics, 10(2):171–178.

Fischbacher, U. and Föllmi-Heusi, F. (2013). Lies in disguise - an experimental study on cheating. *Journal of the European Economic Association*, 11(3):525–547. Publisher: Oxford University Press.

Gerlach, P., Teodorescu, K., and Hertwig, R. (2019). The truth about lies: A meta-analysis on dishonest behavior. *Psychological Bulletin*, 145(1):1–44.

Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95(1):384–394.

Gneezy, U., Kajackaite, A., and Sobel, J. (2018). Lying Aversion and the Size of the Lie. *American Economic Review*, 108(2):419–453.

Hamman, J. R., Loewenstein, G., and Weber, R. A. (2010). Self-Interest through Delegation: An Additional Rationale for the Principal-Agent Relationship. *American Economic Review*, 100(4):1826–1846.

Kandul, S. and Kirchkamp, O. (2018). Do I care if others lie? Current and future effects when lies can be delegated. *Journal of Behavioral and Experimental Economics*, 74:70–78.

Khalmetski, K. and Sliwka, D. (2019). Disguising Lies: Image Concerns and Partial Lying in Cheating Games. *American Economic Journal: Microeconomics*, 11(4):79–110.

Levine, E. E. and Schweitzer, M. E. (2014). Are liars ethical? On the tension between benevolence and honesty. *Journal of Experimental Social Psychology*, 53:107–117.

Levine, E. E. and Schweitzer, M. E. (2015). Prosocial lies: When deception breeds trust. *Organizational Behavior and Human Decision Processes*, 126:88–106.

Michailidou, G. and Rotondi, V. (2019). I'd lie for you. *European Economic Review*, 118:181–192.

Moshagen, M. and Hilbig, B. E. (2017). The statistical analysis of cheating paradigms. *Behavior Research Methods*, 49(2):724–732.

Murphy, R. O., Ackermann, K. A., and Handgraaf, M. (2011). Measuring Social Value Orientation. *SSRN Electronic Journal*.

Palfrey, T. R. and Wang, S. W. (2009). On eliciting beliefs in strategic games. *Journal of Economic Behavior & Organization*, 71(2):98–109.

Ryan, R. M. and Deci, E. L. (2000). Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary educational psychology*, 25(1):54–67.

Soraperra, I., Weisel, O., and Ploner, M. (2019). Is the victim Max (Planck) or Moritz? How victim type and social value orientation affect dishonest behavior. *Journal of Behavioral Decision Making*, 32(2):168–178. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/bdm.2104.

Weisel, O. and Shalvi, S. (2015). The collaborative roots of corruption. *Proceedings of the National Academy of Sciences*, 112(34):10651–10656. Publisher: National Acad Sciences.

A Translated Instructions: Fee and NoFee Conditions

Welcome,

For showing up on time, you will receive 3 euros, which will be paid to you at the end of the experiment. If you have any doubts during the experiment, please address a staff member by raising your hand. If you are using the computer for activities not closely related to the experiment, you will be excluded from the experiment and from any. payment.

Phases

The experiment consists of 3 phases. In each phase, you will be able to earn a sum of money which will depend on your choices and/or the choices of other participants. Each phase will add to the gain in the other phases to define the total gain in the experiment. You will now receive instructions for phase 1. Instructions for phases 2 and 3 will be distributed afterward.

Phase 1: Prospects

Roles and Rounds

At the beginning of this phase, you will be randomly assigned to the role of Subject 1 or Subject 2. You will keep this role for the duration of the phase. In total, you will face 3 rounds of choice in this phase, similar in structure but different in content. In each of the 3 rounds of choice, Subject 1 and Subject 2 will be associated. In each of these rounds, the association between Subject 1 and 2 will change so that no Subject 1 and no Subject 2 will be matched with each other more than once. The identity of the matched subjects will never be revealed.

Decision

In each round, Subject 1 will have to choose between two options: Prospect A and Prospect B.

- **Prospect A** offers the following earnings based on the outcome of the roll of a die:

  - if the result is 2, 4, or 6 (even number), Subject 1 will earn 7 Euros
  - if the result is 1, 3, or 5 (odd number), Subject 1 will earn 3 Euros

The outcome of the lottery is defined by the die thrown by Subject 1. The roll takes place privately and the the result of the roll will be seen only by Subject 1. Subject 1 will report the outcome of the roll and the corresponding earnings will be calculated by the computer.

- Prospect B offers the following earnings based on the outcome of the roll of a die:
  - if the result is 2, 4, or 6 (even number), Subject 1 will earn H Euros
  - if the outcome is 1, 3, or 5 (odd number), Subject 1 will earn L Euro

The values of H and L will change over the 3 rounds and will be specified directly on the monitor. Unlike Prospect A, the outcome is defined by the die tossed by Subject 2. The roll takes place privately, and the result of the roll will only be seen by Subject 2. Subject 2 will report the outcome of the roll, and the computer will calculate the corresponding earnings.

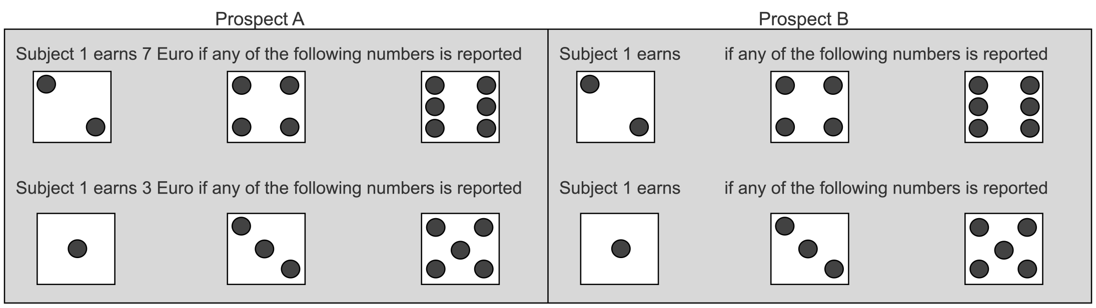

Subject 2 will be asked to report the outcome of the launch before knowing the choice between Prospect A and Prospect B made by Subject 1. If Subject 1 chooses Prospect B, the reported outcome will be relevant for the definition of Subject 1's earnings. If Subject 1 chooses Prospect A, the reported result will not be relevant for the definition of the earnings of Subject 1 and will not be disclosed.

[begin of FEE treatment only] *In each round, before Subject 1 chooses between Prospect A and Prospect B, Subject 2 may request a sum from Subject 1 to report the outcome of the roll of the die. This sum must be between 0 and 1 Euro, in intervals of 10 cents. The amount* *requested will be communicated to Subject 1, and if he/she chooses Prospect B, the requested amount will be subtracted from the lottery profit and paid to Subject* [end of FEE treatment only]

During the phase, Subject 1 will not know the outcome of the roll associated with the Prospect B, and Subject 2 will not be aware of Subject 1's choice between the two prospects.

Earnings

The payoff of Subject 1 in this phase will be given by the outcome of the selected prospect in one of the three rounds chosen at random. Subject 1 will know the gain in this phase only at the end of the experiment. Subject 2's gain in this phase is given by a predefined sum that will be revealed only at the conclusion of the experiment. This amount will not depend in any way on the choices in the experiment.

Phase 2: Expectations

Before knowing the actual choice of the other, both Subject 1 and Subject 2 will be asked to express an expectation about the choice of the other. In total 3 expectations will be expressed that refer to the 3 rounds of Phase 1. More details are provided on the monitor.

Subject 1

Subject 1 will be asked if Subject 2 has reported an odd or even number. The choices are made by assigning a probability to the even (odd) event. Based on the assigned probabilities and upon the actual choice of Subject 2, a sum in EURO is earned. The table below represents the pattern of choice and possible earnings.

Probability that the other reported an EVEN number	0%	25%	50%	75%	100%
Probability that the other reported an ODD number	100%	75%	50%	25%	0%
	○	○	○	○	○
Earnings if the reported number is EVEN	0.00	0.90	1.50	1.90	2.00
Earnings if the reported number is ODD	2.00	1.90	1.50	0.90	0.00

For example,

- A subject assigns 100% probability to the even outcome:
  - if the result reported is even, the gain is 2 EURO;
  - if the result reported is odd, the gain is 0 EURO.
- A subject assigns a 50% probability to the even outcome:
  - if the reported result is equal, the gain is 1.50 EURO;
  - if the result reported is odd, the gain is 1.50 EURO.
- A subject assigns 0% probability to the even outcome:
  - if the result reported is even, the gain is 0 EURO;
  - if the result reported is odd, the gain is 2.00 EURO.

Subject 2

Subject 2 will be asked if Subject 1 has chosen Prospect A or B. The methods of defining earnings are similar to those of Subject 1. The table below provides a representation of the choice setting and possible earnings. For example,

- One subject assigns 100% probability to the choice of Prospect A:
  - if Prospect A has been chosen, the gain is 2 EURO
  - if Prospect B has been chosen, the gain is 0 EURO.
- One subject assigns 50% probability to the choice of Prospect A:
  - if Prospect A has been chosen, the gain is 1.50 EURO
  - if Prospect B has been chosen, the gain is 1.50 EURO.
- One subject assigns 0% probability to the choice of Prospect A:

– if Prospect A has been chosen, the gain is 0 EURO

– if Prospect B has been chosen, the gain is 2.00 EURO.

Earnings

One of the three expectations reported in this phase will be randomly extracted, and the corresponding sum will be paid at the end of the experiment. An expectation will be extracted that refers to a round other than the one randomly extracted for the payment of Phase 1 of the experiment. For example, if round 3 is drawn for the payment of Phase 1, the expectation that refers to round 2 or round 1 will be paid at this stage, randomly.

Phase 3: Allocation

In this stage, you will use ECUs to measure your gains. At the end of the experiment, each ECU will be converted into 2 euro cents. You will make a series of decisions about the distribution of ECUs between you and another participant. The participant will be chosen at random from those in the room. In each of the screens that will be proposed to you, indicate the distribution you prefer by clicking on the corresponding button. There are no right or wrong answers; just make the choice you prefer. The example below shows some possible allocations. In the specific example, the sum of the decision-maker does not change between the various options while that of the other decreases by moving with the choice to the right.

You receive	340	340	340	340	340	340	340	340	340
	○	○	○	○	○	○	○	○	○
The other receives	340	305	270	235	200	165	130	95	60

Earnings

After every participant has made their choices, you will be randomly assigned to the role of Decision Maker or Receiver. If you are a Decision Maker, then one of your choices (selected a case) will define your earnings and that of the other. If you are a Receiver, your earnings will be defined by the choices of another participating Decision Maker.

B Translated Instructions: CPU Condition

Welcome,

For showing up on time, you will receive 3 euros, which will be paid to you at the end of the experiment. If you have any doubts during the experiment, please address a staff member by raising your hand. If you are using the computer for activities not closely related to the experiment, you will be excluded from the experiment and from any. payment.

Phases

The experiment consists of 3 phases. In each phase you will be able to earn a sum of money which will depend on your choices and / or on a random outcome defined by the computer (CPU). Each phase will add to the gain in the other phases to define the total gain in the experiment. You will now receive instructions for phase 1. Instructions for phases 2 and 3 will be distributed afterward

Phase 1: Prospects

Roles and Rounds

In total, you will face 3 rounds of choice in this phase, similar in structure but different in content. In each of the 3 rounds of choice, you will be Subject 1 and you will be matched with a computer (CPU).

Decision

In each round, Subject 1 will have to choose between two options: Prospect A and Prospect B.

- **Prospect A** offers the following earnings based on the outcome of the roll of a die:
  - if the result is 2, 4, or 6 (even number), Subject 1 will earn 7 Euros
  - if the result is 1, 3, or 5 (odd number), Subject 1 will earn 3 Euros

The outcome of the lottery is defined by the die thrown by Subject 1. The roll takes place privately and the the result of the roll will be seen only by Subject 1. Subject 1 will report the outcome of the roll and the corresponding earnings will be calculated by the computer.

- **Prospect B** offers the following earnings based on the outcome of a virtual die roll by the CPU:
  - if the result is 2, 4, or 6 (even number), Subject 1 will earn H Euros
  - if the outcome is 1, 3, or 5 (odd number) Subject 1 will earn L Euro

The values of H and L will change over the 3 rounds and will be specified directly on the monitor. Unlike Prospect A, the outcome is defined by the virtual die thrown by the CPU. Each number always has the same probability of being drawn (= 1/6).

If Subject 1 chooses Prospect B, the reported outcome will be relevant for the definition of Subject 1's earnings. If Subject 1 chooses Prospect A, the reported result will not be relevant for the definition of the earnings of Subject 1 and will not be disclosed.

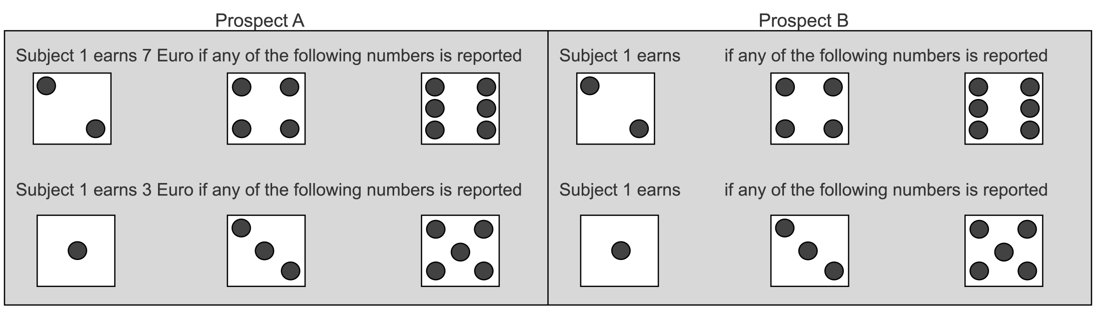

Earnings

The payoff of Subject 1 in this phase will be given by the outcome of the selected prospect in one of the three rounds chosen at random. Subject 1 will know the gain in this phase only at the end of the experiment.

Phase 2: Expectations

Subject 1 will be asked to express an expectation about the choice of CPU. In total 3 expectations will be expressed that refer to the 3 rounds of Phase 1. More details are provided on the monitor.

Subject 1

Subject 1 will be asked if CPU has reported an odd or even number. The choices are made by assigning a probability to the even (odd) event. Based on the assigned probabilities and upon the actual choice of CPU, a sum in EURO is earned. The table below represents the pattern of choice and possible earnings.

Probability that the other reported an EVEN number	0%	25%	50%	75%	100%
Probability that the other reported an ODD number	100%	75%	50%	25%	0%
	○	○	○	○	○
Earnings if the reported number is EVEN	0.00	0.90	1.50	1.90	2.00
Earnings if the reported number is ODD	2.00	1.90	1.50	0.90	0.00

For example,

- A subject assigns 100% probability to the even outcome:
  - if the result reported is even, the gain is 2 EURO;
  - if the result reported is odd, the gain is 0 EURO.
- A subject assigns a 50% probability to the even outcome:
  - if the reported result is equal, the gain is 1.50 EURO;
  - if the result reported is odd, the gain is 1.50 EURO.
- A subject assigns 0% probability to the even outcome:
  - if the result reported is even, the gain is 0 EURO;
  - if the result reported is odd, the gain is 2.00 EURO.

Earnings

One of the three expectations reported in this phase will be randomly extracted, and the corresponding sum will be paid at the end of the experiment. An expectation will be extracted that refers to a round other than the one randomly extracted for the payment of Phase 1 of the experiment. For example, if round 3 is drawn for the payment of Phase 1, the expectation that refers to round 2 or round 1 will be paid at this stage, randomly.

Phase 3: Allocation

In this stage, you will use ECUs to measure your gains. At the end of the experiment, each ECU will be converted into 2 euro cents. You will make a series of decisions about the distribution of ECUs between you and another participant. The participant will be chosen at random from those in the room. In each of the screens that will be proposed to you, indicate the distribution you prefer by clicking on the corresponding button. There are no right or wrong answers; just make the choice you prefer. The example below shows some possible allocations. In the specific example, the sum of the decision-maker does not change between the various options while that of the other decreases by moving with the choice to the right.

You receive	340	340	340	340	340	340	340	340	340
	○	○	○	○	○	○	○	○	○
The other receives	340	305	270	235	200	165	130	95	60

Earnings

After every participant has made their choices, you will be randomly assigned to the role of Decision Maker or Receiver. If you are a Decision Maker, then one of your choices (selected a case) will define your earnings and that of the other. If you are a Receiver, your earnings will be defined by the choices of another participating Decision Maker.