Multi-step delegation and the frequency of immoral decisions: Theory and experiment*

Frauke von Bieberstein†, Eberhard Feess‡ and Natalie Packham§

Version: November 24, 2023

Abstract

When people can increase their payoff by violating a moral norm, they may delegate the decision to dilute their perception of responsibility. This may increase the frequency of norm violations if predominantly those people delegate who would have followed the norm. To analyze this issue, we first develop a model with multiple delegation stages, and private information on lying costs and responsibility dilution. The model shows that the dilution effect does not necessarily lead to more immoral decisions. We then perform a large-scale online experiment where subjects can increase their payoff by lying about the outcome of a lottery. We consider treatments with groups of three players and varying delegation possibilities. We find no evidence that delegation increases the overall lying frequency. Participant behavior aligns with a normal distribution for lying costs and a strongly right-skewed distribution for a rather low dilution effect.

JEL: C92, D02, D63, D71

Keywords: Lying, misreporting, delegation, group decisions, dilution of responsibility

*We are grateful to Gökhan Aydogan, Matthias Blonski, Manuel Grieder, Holger Herz, Robert Kirkby, Michael Kurschilgen, Steffen Lippert, George Loewenstein, Joerg Oechsler, Ronald Peeters, and seminar participants at the 2023 Matterhorn Symposium, the 2023 GEABA conference, the 2023 New Zealand Microeconomics Group Meeting, and the Research Seminar at Victoria University of Wellington for helpful comments and discussion.
Von Bieberstein gratefully acknowledges support from the Förderverein IOP (2021).

†University of Bern, Bern, Switzerland. Email: frauke.vonbieberstein@unibe.ch.

‡Corresponding author. Victoria University of Wellington, Pipitea Campus, Lambton Quay, Wellington, New Zealand. Email: eberhard.feess@vuw.ac.nz.

§Berlin School of Economics and Law, Badensche Str. 52, 10825 Berlin, Germany. Email: packham@hwr-berlin.de

1 Introduction

Delegation by (top)managers in companies boosts efficiency by utilizing the superior knowledge of experts down the line (Demski and Sappington, 1987; Aghion and Tirole, 1997), by increasing the speed of decisions (Yukl and Fu, 1999), and by motivating employees (Fehr et al., 2013). At a macroeconomic level, Akcigit et al. (2021) argue that the efficiency of delegation accounts for more than ten percent of the per capita income difference between the United States and India.

As a dark side, however, delegation may also erode moral standards. In a recent representative survey about ethical issues at the workplace with 14,500 U.S. employees, respondents mentioned that superiors explicitly demanded or implicitly called for rule violations (29%), lying (27%), sacrificing safety (9%), and discrimination (3%) (Ivcevic et al. 2020). Infamous effects of delegation along supply chains include the 2012-fire at a Bangladeshi apparel factory of a supplier of Wal-Mart that killed 117 workers, and the 2014-explosion at a factory in Kunshan City, China, that killed 68 workers of a supplier of General Motors (Chan et al., 2020; Huang et al., 2022).

Why may immoral behavior be more frequently observed when principals delegate decisions to agents? Generally speaking, there are (at least) three possible explanations. First, delegating morally questionable decisions to employees down the line, or to decision makers in developing countries as in the examples mentioned above, may simply reduce the probability of detection, and also the (monetary) consequences to the principal in case of detection. Delegation can then be seen as a device to implement financially beneficial immoral decisions at lower expected costs. Second, principals might be tricked by their agents who follow their own agenda instead of the principals' intentions. This is what top managers in the Volkswagen (VW) Dieselgate scandal claimed, when Michael Horn, head of VW Group of America in 2015, and also VW CEO Martin Winterkorn, asserted that "...this was a couple of software engineers who put this in, for whatever reasons" (Mitchell, 2021).[1] If this was true, immoral decisions taken by

[1] In 2015, the U.S. Environmental Protection Agency (EPA) found that VW had violated the Clean Air Act in over 590,000 diesel motor vehicles because the vehicles were equipped with computer software designed to cheat on federal emissions tests. By 2020, the diesel cheating scandal had cost VW 31.3 billion Euros in fines and settlements (Reuters, 2020). In May 2023, Rupert Stadler, the ex-CEO of VW's Audi AG, was the first top manager to confess in court that, while he did not know that vehicles

agents can be seen as an unwarranted side effect of delegation. Third, principals may delegate immoral decisions in order to reduce their own perception of responsibility, something the literature refers to as the *dilution effect* of moral costs (Hamman et al., 2010). Then, delegation is triggered by psychological motives.

In this paper, we focus exclusively on the last channel, and examine whether the dilution of responsibility leads to a higher frequency of immoral decisions. We first develop a behavioral game theoretical model where (im)moral decisions can be delegated multiple times. We then perform a large scale online experiment, in which subjects have three choices: Taking a moral decision, taking an immoral decision that maximizes their and their team members' payoff, and delegating the decision to the next player. In our experiment, we ignore the potential benefits of delegation, which allows us to cleanly identify the (potential) dilution of responsibility, and also its consequences for the final decision. Our paper is the first that analyzes the impact of (multiple) delegation on the frequency of morally questionable decisions both theoretically and experimentally (see the literature review in Section 2).

In our model, n members of a group decide sequentially between an immoral choice that maximizes each group members' payoff, a moral choice that yields a lower payoff, or delegating the decision to the next player in the line. The game ends if one player decides to not delegate. If $n-1$ players delegate, then the last player in the line needs to make the final decision. In any case, all players receive the same payoff.

We assume that the players' preferences differ from neoclassical standard assumptions in two respects: First, the immoral choice yields player-specific intrinsic moral costs of θ_i. Second, if player i delegates and some other player down the line makes the immoral choice, player i still gets the maximum payoff, while their moral costs decrease to $\tau_i\theta_i$, where $\tau_i \leq 1$ captures the dilution of responsibility. The distributions of θ and τ are common knowledge, but the realizations of θ_i and τ_i are private information.

A first intuition might suggest that the probability of taking the immoral choice increases with the number of players in the line. To see this, consider a group of two players and assume that player 1's dilution of responsibility is either very low ($\tau_1 \to 1$)

had been tampered with, he perceived it as possible and acceptable.

or large ($\tau_1 \to 0$). With $\tau_1 \to 1$, we are effectively back to no-delegation. With $\tau_1 \to 0$, player 1 will either take the immoral choice (for low values of θ_i) or delegate the decision to player 2. They will never take the moral choice, as this is strictly dominated by delegating the decision to player 2 for all θ. Since player 2 is then in the same situation as player 1 without delegation, delegation ultimately leads to more immoral decisions. As delegation thus yields more immoral choices for $\tau_1 \to 0$ and is meaningless for $\tau_1 \to 1$, a sufficient condition for a detrimental effect of delegation is that the probability of taking the immoral decision is monotone in τ_1.

Our model, however, reveals that there is a countervailing effect, and that the impact of τ may well be non-monotone. To see this, consider a player 1 with rather low moral costs θ_1 who would have taken the immoral decision without delegation possibility, but who prefers to delegate if τ_1 is sufficiently low. This leads ceteris paribus to fewer immoral decisions, and the strength of this effect depends on the distributions of θ and τ.

To develop an intuition for the countervailing effects at work and their relative size, we consider several Beta distributions for moral costs θ. Most notably, we find two sufficient conditions for a higher probability of immoral decisions with delegation. The first condition is that the distribution of θ is symmetric and that, without delegation, half of players each make the moral and the immoral choice. The second condition is that θ is uniformly distributed. Both conditions are sufficient for all distributions of τ. There also exist cases, although less obvious, where delegation leads to a lower probability of immoral decisions.

In our experiment, we use the canonical approach where subjects can lie about the outcome of a lottery (Fischbacher and Föllmi-Heusi, 2013; Gneezy et al., 2018; Abeler et al., 2019) as immoral decision. We form groups of three subjects each who observe the same outcome of a binary lottery, which either yields the winner prize W or nothing. As the payoff depends solely on the report, subjects can increase their payoff by reporting WON when they actually LOST the lottery. For similar situations, the literature has pointed to a large degree of heterogenous behavior, with some subjects telling the truth, some reporting the payoff-maximizing outcome, and some lying only moderately in case of non-binary lotteries (see the meta-study by Abeler et al. (2019), and Schudy et al. (2023) for a measurement of individual lying costs). The novelty of our experimental design is that we consider three treatments, which differ with respect to the delegation possibility: In the *no-delegation treatment ND*, one randomly chosen player 1 needs to make a report, and their report determines the group's payoff. In *one-step delegation D1*, player 1 may delegate to player 2, who then needs to make a report. Finally, in *two-step delegation D2*, player 2 may delegate to player 3. Comparing these three treatments allows us to identify if and how delegation itself, and the number of stages, influences the overall degree of lying.

Each subject in our online experiment on Amazon MTurk was randomly assigned to just one treatment and one role in a group (between-subject design). Overall, we had 3,602 subjects who lost the lottery and could misreport the outcome. In line with previous research (Hamman et al., 2010; Bartling and Fischbacher, 2012; Gawn and Innes, 2019a), we find clear evidence for the dilution effect, as 22.5% of all subjects who can delegate do so. This can hardly be explained without the dilution of moral costs, as delegation inevitably reduces the probability for the preferred final outcome. Hence, behavioral models based on lying costs only cannot rationalize this finding. Most importantly, however, delegation opportunities have no statistically significant impact on the overall frequency of lying, which is 55.1% in ND, 55.0% in D1, and 53.1% in D2.

We estimate which distributions of lying costs and of the dilution of responsibility best explain our data, thereby restricting attention to Beta distributions. We find a large heterogeneity of lying costs close to a normal distribution, while the distribution for the dilution effect is largely right-skewed, with an average in the reduction of lying costs of less that 10%.

Summing up, our paper is the first theoretical and experimental analysis of the impact of (multiple) delegation opportunities on the frequency of immoral decisions. Our model follows the literature on (im)moral decisions by assuming that moral costs differ across people and are private information. The dilution effect is captured by a reduction in moral costs when the outcome is reached via delegation. Our model shows that the existence of a dilution effect does not necessarily lead to more final immoral decisions, and we provide a clear intuition for the countervailing effects at work.

In our setting, the downside of delegation from the principal's point of view is that agents might not behave in line with the principal's preferences. In the previous literature, this downside is deliberately neglected: In Hamman et al. (2010), depending on their preferences, principals can choose agents as delegates who are known to implement more or less selfish allocations. Hence, there is little or no uncertainty about the behavior of delegates. In Gawn and Innes (2019a), principals know that delegation yields the same (expected) financial payoff as taking the immoral decision themselves, and the degree of uncertainty is also the same with and without delegation. As a consequence, both papers allow for a clean identification of the existence of the dilution effect, but keep silent about whether delegation opportunities yield more immoral decisions when agents might follow their own preferences. Calibrating the distribution functions that fit our data best suggests that the dilution effect is rather small.

Section 2 relates to the literature. Section 3 presents the model. The experiment is described in Section 4. Results are in Section 5. We conclude and point to further research in Section 6.

2 Related literature

Offering subjects the possibility to lie about the outcome of a lottery has become a canonical experimental approach for analyzing peoples' willingness to increase their payoffs by making morally questionable decisions (see the meta-analyses in Jacobsen et al. (2017) and Abeler et al. (2019)). Robust insights include that many people lie only moderately (Fischbacher and Föllmi-Heusi, 2013), often respond to incentives (Kajackaite and Gneezy, 2017; Charness et al., 2019), care about their reputation (Gneezy et al., 2018; Feess and Kerzenmacher, 2018) and consequences for others (Gneezy, 2005; Erat and Gneezy, 2012), and are influenced by what they perceive as the prevalent social norm (Krupka and Weber, 2013; Bicchieri et al., 2019). Schudy et al. (2023) develop a novel experimental approach for estimating lying costs at an individual level. One stream of the literature finds that people are more inclined to lie as members of a group than when they decide alone (Conrads et al., 2013; Muehlheusser et al., 2015; Kocher et al., 2018). However, keeping financial incentives constant, the lying frequency neither depends on the lying behavior of other group members (Feess et al., 2022) nor of competitors (Dato et al., 2019).

Conversely to the literature just discussed, most papers on delegation with morally questionable decisions use variants of the dictator game. In these games, a principal or their agent can choose between a selfish allocation that benefits the two of them at the expense of a third party, and a (more or less) equal division of the payoff among all subjects (Hamman et al., 2010; Bartling and Fischbacher, 2012; Gawn and Innes, 2019a). In the seminal paper by Hamman et al. (2010), the agent's payoff is independent of the final allocation, but they get an additional amount if the principal selects them as delegates. The delegation opportunity leads to more selfish allocations, because agents who implement these allocations are more often selected. In Gawn and Innes (2019a), dictators also choose between a more or less selfish allocation. The experimental design ensures that the probability that the selfish allocation is implemented is identical when it is chosen directly by the dictator (as a principal) and when it is delegated to an agent. As mainly dictators delegate, who would have otherwise chosen the less selfish allocation, the frequency of selfish allocations is higher with delegation opportunities. Both papers show clearly that a dilution effect exists. We also identify the dilution effect, but find no evidence that this leads to more morally questionable decisions when principals are not informed ex ante about their agents' (average) choices.

Other papers assume that dictators can be punished for allocations that are perceived as unfair. Bartling and Fischbacher (2012) show experimentally that delegating the decision has then the additional advantage that it reduces the punishment frequency. Gurdal et al. (2013) and Oexl and Grossman (2012) also find that principals are less often punished when they delegate the decision.[2] More closely related to our approach are the treatments *without* punishment in Bartling and Fischbacher (2012). For these treatments, however, they have only 35 observations with and 34 without delegation, so that their results are inconclusive: Only 6 subjects delegate, and the overall frequencies

[2]Argenton et al. (2023) show that the reversed effect exists as well, that is, people receive higher rewards when they make fair decisions themselves instead of via delegation.

of selfish allocations are 70% with and 65% without delegation. As the focus in Bartling and Fischbacher (2012) is on punishment, they do not consider the effect of two-step delegation.

A few papers consider delegation opportunities with lying. Erat (2013) uses a variant of the sender-receiver game developed by Erat and Gneezy (2012). The principal in their role as sender can send a true or a wrong message to a receiver, or delegate the decision to an agent. Lying increases the principal's and the agent's payoff at the expense of the receiver. More principals delegate when the lie hurts the receiver to a larger extent, and women delegate more often than men. The effect of delegation is not identified, as there is no treatment without delegation.[3] Gawn and Innes (2019b) compare delegation in a standard dictator game to a dictator game amended by lying. Similar to Erat (2013), lying has a negative impact on a third party. As in Gawn and Innes (2019a), they identify the pure dilution effect by ensuring that the expected payoff is the same with and without delegation. Principals delegate more often in the dictator game with lying, which suggests that delegation also reduces lying costs.

Summing up, our experimental design differs from this literature in three main respects: First, in contrast to Erat (2013) and Gawn and Innes (2019b), there are no negative externalities on other participants, that is, lying affects only the principal and the delegates themselves. Inequity aversion as a potential motive is muted, as all participants get the same payoff in any case.[4] Second, principals do not know the probability with which their delegates take the principals' preferred decision. We hence take one step further: While the preceding literature identifies existence of a dilution effect, we analyze its impact on the overall frequency of immoral decisions. Third, our experiment is the first to consider two-step delegation.

[3] Innes (2022) finds that lying is more frequent in a sender-receiver game when receivers make decisions compared to a situation where these decisions are resembled by computers. One possible explanation is that the receivers' possibility of making an active choice reduces lying costs.

[4] From an applied perspective, both settings seem to be relevant: Gawn and Innes (2019b) argue that considering different allocations is appropriate, as delegation is most meaningful with interpersonal consequences. However, there are also situations such as (delegated) tax fraud or whitewashing of balance sheets where only unrelated people are affected, similar to universities paying liars in experiments.

3 The model

3.1 Setup

There are $n \geq 1$ risk-neutral members of a group who take part in a binary lottery that yields an entitlement to payoff $W \in (0, 1)$ for each group member if the lottery is won and $L = 0$ if the lottery is lost. The lottery is won with probability p. All members of the group observe the same outcome of the lottery.

The n players are randomly ordered into a sequence from 1 to n. If the lottery is won, then all players automatically receive the winner prize W and the game ends. If the lottery is lost, then, when it is their turn, a player reports $r \in (w, l, d)$. The game ends if a player reports either w (for won) or l (for lost), and d means that a player delegates to the next player in the line. If $n - 1$ players delegate, then player n needs to make the final decision by reporting w or l.

If the lottery is lost and player i reports w, they face moral costs of θ_i. If player i delegates and the final report is w, then player i faces moral costs of $\tau_i \theta_i$, where τ_i captures the dilution of moral costs through delegating. Both types are iid, with θ_i having support $[0, 1]$ and $\tau_i \in (0, 1]$, and their distribution functions F and G, respectively, are common knowledge. The realizations of θ_i and τ_i are private information.

Consider first the simplest case with just one player who cannot delegate. This player will report w when the lottery is lost iff $\theta_1 \leq W$ and l else. Define $\widetilde{\theta}$ as the threshold such that this player reports w iff $\theta_1 \leq \widetilde{\theta}$, and note that the same threshold applies to the last player in any sequence with n players. The probability that the player reports w is thus $F(\widetilde{\theta})$, and the probability that they report l is $1 - F(\widetilde{\theta})$.

Consider next player j in a sequence of n players who chooses $r_j \in (w, l, d)$. Denote the probability that the final outcome with n players is W as P_n, and the probability that W is the final outcome if player j delegates by P_{n-j}.

Player j's expected utility is

$$U_j = \begin{cases} 0 & \text{if } r_j = l; \\ W - \theta_j & \text{if } r_j = w; \\ P_{n-j}\left(W - \tau_j \theta_j\right) & \text{if } r_j = d. \end{cases}$$

If player j delegates, their utility is $W - \tau_j \theta_j$ if the final report is w, which happens with probability P_{n-j}. Player j tells the truth if $P_{n-j}\left(W - \tau_j \theta_j\right) \leq 0$. We denote the threshold for lying costs that separates truth-telling from delegation by $\overline{\theta}_j$, that is, player j reports l iff $\theta_j \geq \overline{\theta}_j \equiv \frac{W}{\tau_j}$. As neither the realization of θ nor of τ is known ex ante, this happens with probability $1 - E[F\left(\overline{\theta}_j|\tau_j\right)]$.[5] Furthermore, player j reports w if $W - \theta_j \geq P_{n-j}\left(W - \tau_j \theta_j\right)$, i.e. iff $\theta_j \leq \underline{\theta}_j \equiv \left(\frac{1-P_{n-j}}{1-\tau_j P_{n-j}}\right) W$. This happens with probability $E[F\left(\underline{\theta}_j|\tau_j\right)]$. For lying costs in-between, $\theta_j \in \left(\underline{\theta}_j, \overline{\theta}_j\right)$, player j delegates the decision.

A reduction in the dilution effect (that is, an increase in τ_j) has two consequences: On the one hand, it increases the incentive to report truthfully, as $\frac{\partial \overline{\theta}_j}{\partial \tau_j} = -\frac{W}{\tau_j^2} < 0$. On the other hand, it also increases the incentive to lie, as $\frac{\partial \underline{\theta}_j}{\partial \tau_j} = \frac{(1-P_{n-j})P_{n-j}W}{(1-\tau_j P_{n-j})^2} > 0$. Both conditions express that a reduction in the dilution effect reduces the incentive to delegate.

Note that the critical lying costs for the choice between truth-telling and lying, $\widetilde{\theta}$, which is only taken by the last player in the line, and also the critical lying costs for the choice between truth-telling and delegation, $\overline{\theta}_j$, are both independent of the number of players. The reason is that a player tells the truth iff lying by another player would yield a negative payoff, that is, $W - \tau_j \theta_j < 0$. This condition is independent of the probability P_{n-j} that delegating yields a lie as final outcome, and therefore also independent of the expectations on the behavior of other group members. By contrast, the critical lying costs for the choice between lying and delegating, $\underline{\theta}_j$, depend on P_{n-j} (since $\frac{\partial \underline{\theta}_j}{\partial P_{n-j}} = -\frac{(1-\tau_j)}{(1-\tau_j P_{n-j})^2} W < 0$), as a higher probability that someone down the line lies increases the prospects of getting the money at lower moral costs.

[5]The notation $F(\cdot|\tau_1)$ emphasizes the dependence of the argument on τ_1.

3.2 One-step delegation (two players)

To illustrate the effects at work, we first compare the cases with $n = 1$ and $n = 2$. With just one player, we have $P_1 = F(\widetilde{\theta})$. The dilution effect τ doesn't matter, as player 1 cannot delegate. For $n = 2$, we obtain

$$P_2 = \mathbb{E}[F(\underline{\theta}_1|\tau_1)] + \left[\mathbb{E}[F(\overline{\theta}_1|\tau_1)] - \mathbb{E}[F(\underline{\theta}_1|\tau_1)]\right] \cdot F(\widetilde{\theta}). \tag{1}$$

With probability $\mathbb{E}[F(\underline{\theta}_1|\tau_1)]$, player 1 lies in a sequence with two players. With probability $\mathbb{E}[F(\overline{\theta}_1|\tau_1)] - \mathbb{E}[F(\underline{\theta}_1|\tau_1)]$ player 1 delegates. Then, player 2 is in the same situation as player 1 without delegation (and as any last player in a sequence), and hence lies with probability $F(\widetilde{\theta})$. Comparing P_2 in Equation (1) to $P_1 = F(\widetilde{\theta})$ yields

$$P_2 \geq P_1 \iff \left(\mathbb{E}\left[F(\overline{\theta}_1|\tau_1)\right] - F(\widetilde{\theta})\right) F(\widetilde{\theta}) \geq \left(F(\widetilde{\theta}) - \mathbb{E}\left[F(\underline{\theta}_1|\tau_1)\right]\right)(1 - F(\widetilde{\theta})). \tag{2}$$

We now separate the two effects that may lead to more or less lying with delegation, and refer to $\left(\mathbb{E}\left[F(\overline{\theta}_1|\tau_1)\right] - F(\widetilde{\theta})\right) F(\widetilde{\theta})$ as the partial effect that leads to a gain in the lying probability through delegation. Thereby, $E\left[F(\overline{\theta}_1|\tau_1)\right] - F(\widetilde{\theta})$ is the probability that a player delegates instead of telling the truth, and $F(\widetilde{\theta})$ is the probability that the delegate lies. The LHS therefore captures the gain in the final lying probability from the fact that some players who would have told the truth delegate, and that some of the delegates lie. Analogously, the RHS of Inequality (2) expresses the partial effect that leads to a decline in the lying probability through delegation. $F(\widetilde{\theta}) - E\left[F(\underline{\theta}_1|\tau_1)\right]$ is the probability that a player who would have lied without the possibility of delegation delegates, and $(1 - F(\widetilde{\theta}))$ is the probability that the delegate tells the truth. The RHS therefore captures the decline in lying from the fact that some players who would have lied delegate, and that some of these delegates tell the truth. We find:

Proposition 1. The probability of lying ($r = w$) with $n = 2$ players may be higher than, equal to or lower than with $n = 1$ players.

Proof. We prove by examples. See Figure 1 for details.

The relative size of the two countervailing effects of delegation varies with the distribution of types. There are distributions of θ where delegation leads to a higher probability of lying for *all* distributions of τ. However, there also exist distributions of θ *and* τ that decrease the probability of lying.

Proposition 2. *(i)* If the distribution of θ is symmetric and $\widetilde{\theta}$ is the median, then $P_2 > P_1$, for all distributions of τ_1. *(ii)* If θ is uniformly distributed, then $P_2 > P_1$ for all distributions of τ_1. *(iii)* There exists $\widetilde{\tau}_1 > 0$ such that if τ_1 is restricted to $(0, \widetilde{\tau}_1]$, then $P_2 > P_1$. *(iv)* If $P_2(\tau_1) - P_1$ is monotone or convex as a function of τ_1, then there exists $\widetilde{\tau}_1$ such that $P_2 > P_1$ for any τ_1 with $E[\tau_1] \leq \widetilde{\tau}_1$.

Proof. See Appendix.

Note that $\widetilde{\tau}_1$ in (iii) and (iv) depends on $\widetilde{\theta}$ and on the distribution of θ.

The four figures show Beta distributions of θ, a family of distributions bounded on $[0, 1]$ and characterized by two parameters α and β. They include the uniform ($\alpha = \beta = 1$), symmetric ($\alpha = \beta$), left-skewed ($\alpha > \beta$) and right-skewed ($\alpha < \beta$) distributions as special cases. In the symmetric case, if $\alpha = \beta$ are large enough, then the Beta distribution approximates a Normal distribution. In all figures, we set $\tau = 0.8$, and in the first three figures, we set the winner prize to $W = 0.5$. Recall that, without delegation, types tell the truth iff $\theta \geq \widetilde{\theta} = W$.

The upper left Figure 1(a) shows a symmetric Beta distribution ($\alpha = \beta = 10$), so $W = 0.5$ is the median of the distribution. Therefore, without delegation, the probabilities that a player reports honestly ($r = l$) or lies are identical, $F(\widetilde{\theta}) = 1 - F(\widetilde{\theta}) = 0.5$. Then, delegation yields more lying iff more players 1 switch from truth-telling than from lying to delegation, that is, $P_2 \geq P_1$ iff $E\left[F(\overline{\theta}_1|\tau_1)\right] - F(\widetilde{\theta}) \geq F(\widetilde{\theta}) - \mathbb{E}\left[F\left(\underline{\theta}_1|\tau_1\right)\right]$. The numbers and shaded areas in Figure 1(a) show that the probability that a player 1 switches from truth-telling to delegation is 0.369, compared to only 0.271 for switching from lying to delegation. As half of the delegates lie, this translates directly into more lying with delegation opportunity, as shown by the numbers in brackets. The reason why this extends to all symmetric distributions with $\widetilde{\theta}$ as median (part (i) of Proposition 2) is that the utility with truth-telling is independent of types θ, while the utility from lying is the higher the lower is θ. As delegating thus becomes unattractive when θ_1 decreases, the region $[\widetilde{\theta}, \underline{\theta}_1]$ where people switch from lying to delegating is

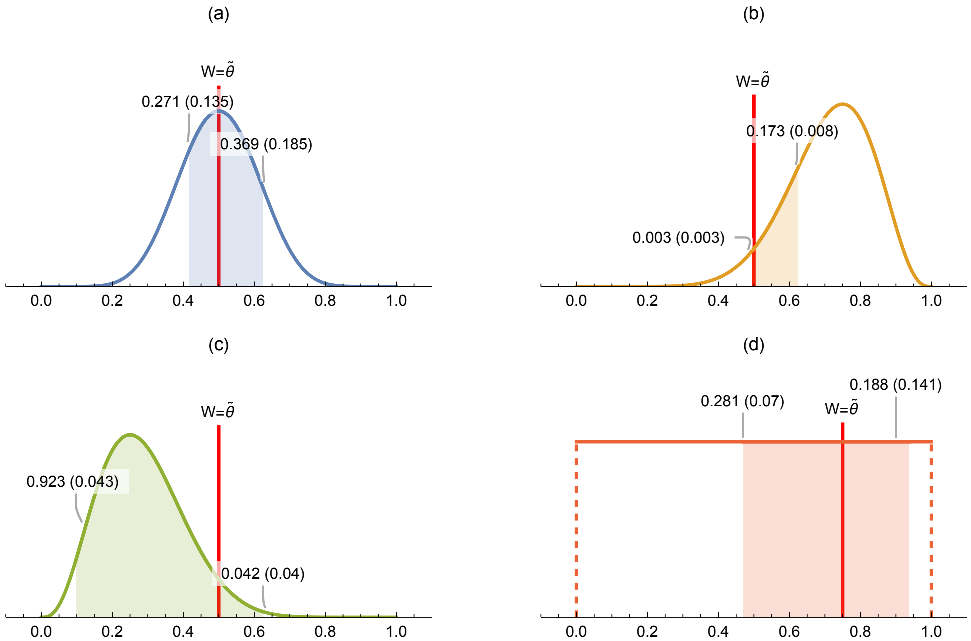

Figure 1: The figures show the probabilities that types delegate instead of telling the truth or lying for different Beta distributions of θ with $\tau = 0.8$. The straight line shows the threshold lying costs $\widetilde{\theta} = W$. The shaded areas and numbers to the right (to the left) of W show the probabilities for player 1 switching from truth-telling (lying) to delegating, and hence the intervals $[\widetilde{\theta}, \overline{\theta}]$ ($[\underline{\theta}, \widetilde{\theta}]$). The numbers in parentheses to the right (to the left) show how delegating instead of truth-telling leads to more (less) lying.

smaller then the region $[\widetilde{\theta}, \overline{\theta}_1]$ where people switch from telling the truth to delegating.

For the right-skewed Beta distribution in Figure 1(b) ($\alpha = 10$, $\beta = 4$), there are two countervailing effects: On the one hand, far more players switch from truth-telling than from lying to delegation (0.173 compared to just 0.003), because there is a high probability mass (close) to the right of $\widetilde{\theta}$, and a low probability mass (close) to the left of $\widetilde{\theta}$. However, the fact that most lying costs are above $\widetilde{\theta}$ also implies that most delegates tell the truth ($1 - F(\widetilde{\theta}) = 0.954$). Therefore, the gain and decline of the lying probability through delegation are both small: the gain because few delegates lie, and the decline because few liars switch to delegation. Still, however, the gain of 0.008 exceeds the decline of 0.003.

For the left-skewed Beta distribution in Figure 1(c) ($\alpha = 4$, $\beta = 10$), the countervailing effects go in the opposite directions: More players switch from lying than from truth-telling to delegation (0.923 vs. 0.042), but the majority of delegates lies ($F(\widetilde{\theta}) = 0.0461$). The gain and decline are again very small, but now the incentives to switch outweigh the behavior of delegates, which yields slightly less lying with delegation (gain of 0.043 vs. decline of 0.04). In any case, the impact of delegation is only moderate due to countervailing effects: Whenever many people switch from truth-telling to delegation, then only few delegates lie, and vice versa.

For uniform distributions ($\alpha = \beta = 1$), delegation leads to more lying even when $\widetilde{\theta}$ is not the median of the distribution (part (ii) of Proposition 2). For our example with $\widetilde{\theta} = W = 0.75$ in Figure 1(d), the probability that those who would have lied delegate is 0.281, compared to only 0.188 for those who would have told the truth. Nevertheless, delegation yields more lying, as 75% of delegates lie. The gain in the lying probability is thus $0.188 \cdot 0.75 = 0.141$, compared to a decline of $0.281 \cdot 0.25 = 0.07$.

Part (iii) of Proposition 2 expresses that, if the maximum τ_1 is sufficiently small, then delegation always leads to more lying, where the exact upper bound for τ_1 depends on the distribution of θ and on $\widetilde{\theta}$. The reason is that, if τ_1 is small, then truth-telling is dominated by delegating the decision, as one might get the winner prize at virtually no moral costs.

Finally, part(iv) of Proposition 2 states that, if the difference in lying with and without delegation is monotone or convex in τ_1, then delegation yields more lying if the expected dilution effect is sufficiently large. In this case, hardly anyone reports honestly, as this is dominated by delegation, thereby getting the winner prize with positive probability at (almost) no cost. By contrast, lying is still attractive for those with low lying costs, as they might not get the prize at all when they delegate.

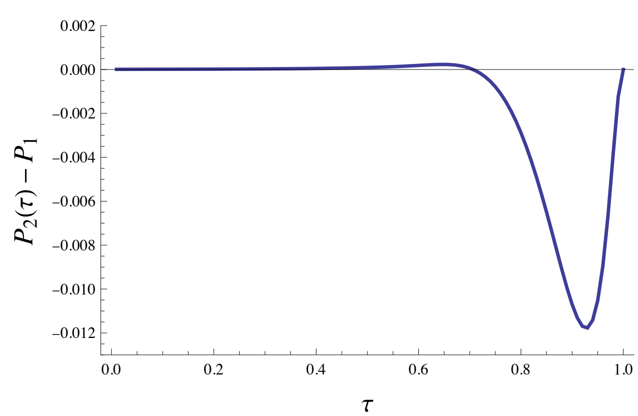

Figure 2: The figure shows $P_2(\tau) - P_1$ with a Beta distribution of θ with $\alpha = 0.4$ and $\beta = 10$ (as in figure 1(c)).

To illustrate the impact of τ, consider again the left-skewed Beta distribution with $\alpha = 4$, $\beta = 10$ from Figure 1(c). In our example with $\tau = 0.8$, we found that delegation leads to a small decrease in the lying probability by 0.003. If we fix $\tau = 0.5$ instead, the lying probability increases by 0.009, as the higher dilution effect makes delegation more attractive for those who would otherwise have told the truth. Figure 2 plots the overall increase in the lying probability, $P_2(\tau) - P_1$, against the dilution effect τ.

For sufficiently small τ, the difference in lying probabilities is positive (see part (iii) of the Proposition), and this holds up to $\tau = 0.706$. Then, lying decreases compared to the case without delegation, and is then identical for $\tau = 1$ (as no player delegates).

3.3 Multi-step delegation (n players)

For the general case with n players, recall first that the critical lying cost for the choice between truth-telling and delegation, $\overline{\theta}_j$, is independent of the number of players. By contrast, the critical lying costs for the choice between lying and delegating, $\underline{\theta}_{j,n}$, depend on the probability for the final outcome, P_{n-j}, and hence also on the number of players and the own position in the line. Lemma 1 connects $\underline{\theta}_{j,n}$ to the probability of lying as final result.

Lemma 1. Suppose the number of remaining players $n - j$ in a sequence increases. Then, the lower bound $\underline{\theta}_{j,n+1}(\tau_j)$ that separates delegation and lying decreases with n iff the probability of lying by subsequent players increases, i.e., $\underline{\theta}_{j,n+1}(\tau_j) \leq \underline{\theta}_{j,n}(\tau_j) \iff P_{n+1-j} \geq P_{n-j}$.

Proof. See Appendix.

To see this, recall that the threshold $\underline{\theta}_{j,n+1}(\tau_j)$ refers only to players who would have lied if they cannot delegate, thus $\underline{\theta}_{j,n+1}(\tau_j) \leq \widetilde{\theta}$. Any player j prefers delegating to lying iff the decrease in the probability of getting the winner prize W is overcompensated by the reduction in lying costs from θ_j to $\tau_j \theta_j$. Therefore, the region in which player j prefers delegation over direct lying increases – which is equivalent to a decrease in $\underline{\theta}_j$ – with the probability of getting the money in case of delegation. We get

Proposition 3. *(i)* If the distribution of θ is symmetric and unimodal, and if $\widetilde{\theta}$ is the median, then $P_{n+1} > P_n$, for all n. *(ii)* If θ is uniformly distributed, then $P_{n+1} > P_n$, for all n. *(iii)* Suppose $P_2 > P_1$. Then a sufficient condition that $P_{n+1} > P_n$ holds for all $n > 1$ is that $F''(\widehat{\theta}) \geq 0$, for all $\widehat{\theta} \leq W$.

All of the statements hold regardless of the distribution of τ.

Proof. See Appendix.

Proposition 3 extends the main insights of Proposition 2 from 2 to n players: Parts (i) and (ii) correspond to parts (i) and (ii) of Propositions 2, where the intuition is again that, under these conditions, the incentive to switch from truth-telling to delegation is higher than the incentive to switch from lying. Furthermore, for any distribution of τ, there exist distributions of θ such that the probability of lying as final outcome increases with the number of players. By an induction argument, it is sufficient to consider the case of player 1 when increasing the number of players from n to $n+1$. Assume $P_n > P_{n-1}$. Again, there are the two countervailing effects discussed for the case with only two players: a gain (decline) in the total probability of lying from players who would tell the truth (who would lie) in an n-player game. As $P_n > P_{n-1}$, the incentive to delegate instead of lying increases, hence $\underline{\theta}_{1,n+1} \leq \underline{\theta}_{1,n}$ (see Lemma 1). The condition $F''(\hat{\theta}) \geq 0$ for all $\hat{\theta} \leq W$ implies that the density of θ is non-decreasing in the region below $\underline{\theta}_{1,n}$ (which is smaller than W). This bounds the increase in the probability of delegating instead of lying, and as a consequence, the effect that people delegate instead of telling the truth dominates.

4 Experimental design and procedure

4.1 Design

In our online experiment, we randomly formed groups of three subjects. Groups took part in a lottery with a winning price of 1.8 USD per subject and a losing price of 0.9 USD per subject.[6] Subjects knew that the probability for the winning prize is 25%, that all of them observe the same outcome of the lottery, and that they get the same prize which depends only on the report. The outcome of the lottery was observable to

[6]For the instructions see the online appendix.

us.[7]

The experiment consisted of three treatments, the *no-delegation treatment ND*, the *one-step delegation treatment D1*, and the *two-step delegation treatment D2*. Each subject performed in just one treatment and one role in a group (between-subject design). In all treatments, subjects were randomly assigned to the role of either player P1, P2, or P3. In ND, we asked all three subjects for their decisions, and informed them that only the report of one player would determine the outcome. Subjects in ND could only report that they won (w) or lost (l) the lottery. In D1, P1 could report w or l, or delegate (d) the decision to P2. If P1 delegated, then P2 needed to make the final decision. P2 was asked for their decision in case P1 delegates. P3 was passive. The only difference of D2 to D1 was that P2 could now also delegate, in which case P3 made the final decision. P2 and P3 were asked for their decisions in case they have to decide. We thus applied the strategy method in all three treatments. Beforehand, we ran a pilot where we used the strategy method on 57 subjects and the direct response method on 60 subjects in treatment ND. Results for both groups were similar.[8]

After the main part of the experiment, subjects were asked for their belief about the percentages of other subjects who report w, l and d (if possible) in their respective roles. For instance, in treatment D1, we asked all subjects to estimate how many out of 100 participants in the role of P1 would choose each of the three reports. Similarly, we asked all subjects to estimate how many out of 100 participants in the role of P2 choose w or l. As the belief elicitation is not at the heart of our research interest, we did not incentivize it.

After the belief question, we measured "need for control" by four items (De Rijk et al., 1998) on a scale from 0 (lowest preference to be in control) to 10.[9] We then

[7]The literature finds that lying is reduced when the outcome is observable to the experimenter (Gneezy et al., 2018; Abeler et al., 2019). This should be no concern for our setting, as we are only interested in treatment (and not in size) effects. Observability saves on costs, as each lie is identifiable. Schudy et al. (2023) provide an innovative experiment to separate between internal lying costs and reputation effects at the individual level.

[8]The number of subjects given in the text counts only people in groups that lost the lottery. Overall, 75 subjects took part in each of the two versions of the ND treatment, with and without strategy method. In the treatment with strategy method, 56% of subjects lied, and without the strategy method, 60% lied ($p = 0.672$; χ^2-test).

[9]The four items are "I prefer giving orders instead of receiving them.", "I prefer having control over what I do and the way I do it.", "I prefer doing my own planning.", and "I prefer being able to set the pace of my tasks.". The reliability of the measure is good (Cronbach's alpha$\,=\,$0.803).

asked for the preference to be a delegate by the question "Are you generally a person who likes that others delegate decisions to you?" on a scale from 0 (lowest preference to be a delegate) to 10. This preference could affect the behavior of subjects in the role of a delegate. In order to assess the perceived responsibility for own decisions, we asked subjects the two questions "If I make a decision and lie to another person, I am responsible for this lie" and "If I make a decision that harms a third person, I am responsible for this harm" on a scale from 0 (lowest responsibility) to 10. We assessed the perceived responsibility for the decisions of the delegate with the two questions "If I delegate a decision and the delegate lies to another person, I am responsible for this lie" and "If I delegate a decision and the delegate's decision harms a third person, I am responsible for this harm" on the same scale. We view this as a (rough) proxy for the dilution effect. We measured the willingness to take risks by the question "Are you generally a person who is willing to take risks, or do you try to avoid taking risks?" on a scale from 0 (lowest willingness to take risks) to 10 (Dohmen et al., 2011). Finally, we asked for gender and year of birth.

4.2 Procedure

We preregistered our study in the AEA RCT Registry (DOI AEARCTR-0008099). Our experiment was run on Amazon MTurk, which is a large online platform where people can participate in research and business studies, between June and August, 2022. We announced a scientific study and a survey on individual decision-making. To ensure high data quality, we required subjects to be fluent in English, to reside in the USA, to be at least 18 years old, to have at least 500 approved HITs, and to have a HIT approval rate of at least 95%. All subjects were allowed to participate just once. We implemented safeguards to avoid that subjects could restart the survey. We informed them that the study takes about 10 minutes, involves filling out a short survey, and making a decision. Subjects also knew that the payment would be 1.00 – 1.90 USD. If subjects were interested in participating, they followed a link taking them to the first page of our study (hosted on Heroku). This first page was a consent form, and only subjects giving their consent entered the study. The experiment was computerized using oTree (Chen et al., 2016).

We conducted a pilot for treatment ND, in which about 60% of participants lied. Based on this data, a two-sided χ^2-test, an error probability of 0.05, and a power of 0.80, we required about 600 group decisions per treatment to detect differences in the lying frequency per treatment of 7.5 percentage points. Thus, we aimed to have 600 participants losing the lottery for treatment ND (where all three players are in the same role), and $600 \cdot 3 = 1,800$ participants losing the lottery for each of the two delegation treatments. Overall, we thus required $4,200$ participants losing the lottery.

$7,074$ subjects participated in the study. We wanted to make sure that the subjects read the instructions carefully. Therefore, we asked three comprehension questions, which referred to the group composition and to the prize structure. Subjects knew that they could not proceed when they made more than one mistake in answering these questions. Overall, 65.5% (14.0%) of subjects made no mistake (just one mistake), so $5,628$ subjects proceeded to the experiment. 20.4% of subjects needed to be excluded because they made more than one mistake.[10] We also included an attention check by asking subjects to select a particular choice during the final survey. In total, 94.4% of subjects made the correct choice. We control for this attention check in the regressions. Of the $5,628$ subjects, $4,200$ (74.6%) subjects lost the lottery, with 615 of them in ND, 598 groups in D1, and 597 groups in D2. Dropping subjects P3 in D1 who did not make a decision leaves us with the final sample of $3,602$ subjects. The mean age is 38.5 years (SD = 11.6).

On average, treatments ND, D1, and D2 took subjects 5, 5.5, and 6 minutes, respectively. Including the show-up fee of 0.10 USD, subjects earned on average 1.49 USD, which translates into about 14.90 USD per hour.

[10] These subjects earned only the show-up fee.

5 Results

5.1 The overall impact of delegation possibilities

Recall that we can observe which groups lost the lottery, so we consider only these cases. Table 1 shows the subjects' decisions by treatments and their position in the game. The three blocks ND (*No Delegation*), D1 (*One-Step Delegation*), and D2 (*Two-Step Delegation*) display the behavior of subjects in the respective treatments, separated by their position in the game. For calculating the first percentage ("all"), the denominator consists of all subjects in a treatment. For calculating the second percentage ("no del"), the denominator consists only of those subjects in a treatment who did not delegate (and thus reported w or l).

Table 1. Overview of decisions by treatments and positions

		ND			D1			D2		
		n	% (all)	% (no del)	n	% (all)	% (no del)	n	% (all)	% (no del)
P1	Truth	276	44.9%	44.9%	206	34.5%	45.8%	215	36.0%	48.1%
	Lie	339	55.1%	55.1%	244	40.8%	54.2%	232	38.9%	51.9%
	Del	-	-	-	148	24.8%	-	150	25.1%	-
P2	Truth	-	-	-	254	42.5%	42.5%	211	35.3%	43.0%
	Lie	-	-	-	344	57.5%	57.5%	280	46.9%	57.0%
	Del	-	-	-	-	-	-	106	17.8%	-
P3	Truth	-	-	-	-	-	-	272	45.6%	45.6%
	Lie	-	-	-	-	-	-	325	54.4%	54.4%
	Del	-	-	-	-	-	-	-	-	-
Sum		615	-	-	1,196	-	-	1,791	-	-
FREQ	Truth	-	44.9%	-	-	45.0%	-	-	46.9%	-
	Lie	-	55.1%	-	-	55.0%	-	-	53.1%	-

Notes: The table shows, for all roles in the game, the absolute numbers and frequencies of subjects who tell the truth, lie, and delegate. The treatments are *No Delegation* (ND), *One-Step Delegation* (D1), and *Two-Step Delegation* (D2). P1, P2 and P3 refer to the roles of player 1, player 2, and player 3, respectively. For all decisions, we apply the strategy method. By definition of the treatments, P3s make no decision in ND and D1, and P2s make no decision in ND. The first percentage ("all") shows the subjects that take the respective choice. The second percentage ("no del") is the percentage of subjects who do *not* delegate and report W or L. The final line FREQ shows the *expected* overall lying and truth telling frequencies, that is, the decisions by P2s and P3s are multiplied with the probabilities that their respective predecessor delegates the decision.

Most importantly, Table 1 shows the expected overall lying frequency (FREQ), where the decisions by P2 and P3 are multiplied with the probabilities that their respective predecessor delegated the decision. This lying frequency hardly differs across treatments, with a minimum of 53.1% for D2, and a maximum of 55.1% for ND. As the differences are not significant, our data reject the hypothesis that delegation leads to a higher overall lying frequency (ND vs. D1: $p = 0.971$; ND vs. D2: $p = 0.480$; D1 vs. D2: $p = 0.506$; χ^2-tests). Note that we cannot run a regression with the final outcome as dependent variable, because the expected lying frequencies are calculated at the group level, while controls are at the individual level.

Table 2 restricts attention to subjects who either tell the truth or lie; regardless of whether they decided against delegation, or could not have delegated because they were the last players.

Table 2. Honest and dishonest reports by treatments

		Truth	Lie
ND	n	276	339
		(44.9%)	(55.1%)
D1	n	460	588
		(43.9%)	(56.1%)
D2	n	698	837
		(45.5%)	(54.5%)

Notes: The table reports the decisions of subjects who tell the truth or lie. Percentages are given in parentheses.

In line with the expected overall lying frequency shown in Table 1, we find no treatment effect for the decisions of subjects who did not delegate. We observe the highest lying frequency of 56.1% in D1, and the lowest with 54.5% in D2. All differences across treatments are not significant (ND vs. D1: $p = 0.696$; ND vs. D2: $p = 0.802$; D1 vs. D2: $p = 0.428$; χ^2-tests). The very small effect sizes and narrow confidence intervals support this null result (ND vs. D1: Cohen's $d = -0.02$, 95% CI $[-0.12, 0.08]$; ND vs. D2: $d = 0.01$, 95% CI $[-0.08, 0.11]$; D1 vs. D2; $d = 0.03$, 95% CI $[-0.05, 0.11]$).

As decisions are now at the individual- instead of the group-level, we can test if our null result is robust to adding controls. Table 3 shows a linear probability model (OLS) where the dependent variable takes the value 0 (1) if the decision maker tells the truth (lies).[11]

[11] All results are qualitatively the same for logit and probit models.

As Table 2, the model includes only subjects who did not delegate.[12] The main independent variable is the treatment, with no delegation (ND) as reference category.

Table 3. Regression analysis on final decisions

	Decision maker lies instead of telling the truth		
	(1)	(2)	(3)
D1	0.006	0.016	-0.012
	(0.026)	(0.027)	(0.025)
D2	-0.005	0.007	-0.009
	(0.024)	(0.027)	(0.024)
Male	0.010	-0.005	-0.005
	(0.018)	(0.018)	(0.016)
Other gender	-0.127	-0.081	-0.112
	(0.098)	(0.097)	(0.084)
Age	-0.004***	-0.003***	-0.001*
	(0.001)	(0.001)	(0.001)
Risk		0.016***	0.014***
		(0.004)	(0.004)
Need for control		0.009	0.006
		(0.006)	(0.006)
Likes to be delegate		0.004	-0.002
		(0.004)	(0.004)
Responsibility own		-0.032***	-0.028***
		(0.006)	(0.005)
Responsibility del		0.009**	0.004
		(0.004)	(0.004)
Attention check		-0.062	-0.068*
		(0.038)	(0.037)
Final decision maker		0.012	0.009
		(0.020)	(0.018)
Belief FREQ lie			0.867***
			(0.028)
Constant	0.686***	0.695***	0.235***
	(0.037)	(0.072)	(0.069)
Observations	3,025	3,025	3,025

Notes: OLS regression on the final decision. Robust standard errors in parentheses. The dependent variable takes the value 0 (1) if the final decision maker tells the truth (lies). The main independent variable of interest is the treatment dummy with *No Delegation* (ND) as reference category. For gender, female is reference category, and other gender includes "non-binary or other" and "prefer not to answer", with 21 and 10 observations, respectively. Age is measured in years. Risk is measured by the question "Are you generally a person who is willing to take risks, or do you try to avoid taking risks?" on a scale from 0 (lowest willingness to take risks) to 10. Likes to be delegate is measured by the question "Are you generally a person who likes that others delegate decisions to you?" on a scale from 0 (lowest preference to be a delegate) to 10. Attention check is a binary variable that takes the value 0 (1) if the answer to the attention check was correct (wrong). Final decision maker is a binary variable that takes the value 0 (1) if the decision maker had (not) the option to delegate. Our measures for need for control, perceptions of responsibility, and belief FREQ lie are described in Appendix B. *, ** and *** document significance at the 10-, 5- and 1-percent level, respectively.

[12] The regressions include 3,025 instead of the 3,198 participants in Table 2, as the survey on personal attributes did not work in one experimental session.

In line with the χ^2-test, we find that the treatments have no significant impact on the lying probability. Column (1) includes only the treatment dummies, gender, and age. Gender is insignificant, and age significantly negative throughout. Interactions with gender and age are not significant in any specification and therefore omitted.

The personal controls in column (2) go in the expected directions: The partial correlation between the willingness to take risks and lying is positive, and a higher feeling of responsibility for the own behavior correlates negatively with lying. Less straightforward, a higher feeling of responsibility for the behavior of delegates leads to more lying compared to truth-telling. In line with other studies on morally questionable decisions, column (3) shows that the belief about the lying frequency of other subjects is positively correlated with lying. This needs to be interpreted with caution, as subjects might rationalize their own lying behavior by stating that lying is common, thereby constituting a (descriptive) social norm.

5.2 The impact of the position in the game

Next, we analyze if the position in the game influences the decision. Table 4 shows all subjects who could delegate, that is, P1s in treatment D1, and P1s and P2s in treatment D2.[13]

Table 4. Decisions by roles and treatments

		Truth	Lie	Del
P1 in D1	n	206	244	148
	in %	34.5%	40.8%	24.8%
P1 in D2	n	215	232	150
	in %	36.0%	38.9%	25.1%
P2 in D2	n	211	280	106
	in %	35.3%	46.9%	17.8%
Total	n	632	756	404
	in %	42.2%	35.3%	22.5%

Notes: The table reports the decisions of subjects who had the possibility to delegate.

Table 4 shows that the percentage of subjects who delegate is basically the same

[13]This information is already included in Table 1, and summarized in Table 4 only for ease of inspection.

for P1s in D1 and in D2 with 24.8% and 25.1%, respectively, but considerably lower for P2s in D2 with 17.8%. In total, 22.5% of our subjects delegate if this is possible, supporting the existence of a dilution effect.

Interestingly, the lower delegation frequency for P2s in D2 leads solely to a higher lying frequency, while the percentage of subjects telling the truth is virtually the same in all treatments. While we can only speculate about the reasons, we opine that the following two, possibly complementary, explanations are plausible: First, as P2s in D2 are the only subjects who are delegates themselves, they may interpret the delegation decision of P1s as an invitation to make the final decision, and therefore delegate less often. Second, P2s in D2 may interpret P1s decision to delegate as a signal on the acceptability of lying, which may then lead to more lying compared to truth-telling. In any case, the difference between P1s in D1 and P1s in D2 is not significant with $p = 0.776$, while the comparison of P2s in D2 to P1s in D1, and also to P1s in D2, is significant, with $p = 0.009$ and $p = 0.002$, respectively (χ^2-tests). The reason why this difference does not translate into an overall higher lying frequency in D2 compared to D1 is that P1s and P3s in D2 tell the truth somewhat more often compared to P1s and P2s in D1 (see Table 1).[14]

The corresponding Table 5 shows results from a multinominal logit. The decision as dependent variable takes the value 0 for lying, 1 for truth-telling, and 2 for delegating. As in Table 4, we consider only subjects who can choose among all three possibilities, that is, all subjects in ND, P2s in D1, and P3s in D2 are excluded.[15] The position in the game is the main independent variable of interest, and takes the value 0 for P1s in D1, 1 for P1s in D2, and 2 for P2s in D2.

The results from the multinominal logit support those from the χ^2-test. Recalling that the reference category for the dependent variable is lying, the first (last) three columns in Table 5 show the impact of the independent variables on the probability to tell the truth (to delegate) instead of lying. Columns (1) and (4) include only the position in the game, gender, and age. Columns (2) and (5) add our personality

[14]The differences are not statistically significant; P1s in D1 vs. P2s in D2: $p = 0.776$ as shown in the text, and P2s in D1 vs. P3s in D2: $p = 0.283$ (χ^2-tests).

[15]Regressions are based on $1,673$ participants, because the survey on personal attributes did not work in one experimental session.

Table 5. Multinominal logit on all three decisions

	Decision maker tells the truth instead of lying			Decision maker delegates instead of lying		
	(1)	(2)	(3)	(4)	(5)	(6)
P1 in D2	0.110	0.093	0.097	0.061	0.057	0.067
	(0.139)	(0.142)	(0.155)	(0.153)	(0.158)	(0.163)
P2 in D2	-0.101	-0.105	-0.021	-0.497***	-0.493***	-0.434**
	(0.137)	(0.140)	(0.155)	(0.160)	(0.163)	(0.169)
Male	-0.031	0.033	0.058	-0.168	-0.081	-0.055
	(0.113)	(0.116)	(0.127)	(0.129)	(0.134)	(0.139)
Other gender	0.239	-0.074	0.315	-0.419	-0.891	-0.572
	(0.578)	(0.578)	(0.633)	(0.824)	(0.838)	(0.915)
Age	0.016***	0.011**	0.008	0.017***	0.012**	0.011*
	(0.005)	(0.005)	(0.006)	(0.005)	(0.006)	(0.006)
Risk		-0.091***	-0.086***		-0.057*	-0.054
		(0.027)	(0.029)		(0.032)	(0.033)
Need for control		-0.001	0.012		-0.098**	-0.089*
		(0.042)	(0.046)		(0.047)	(0.049)
Likes to be delegate		-0.047*	-0.025		-0.128***	-0.114***
		(0.026)	(0.029)		(0.030)	(0.031)
Responsibility own		0.166***	0.129***		0.180***	0.158***
		(0.037)	(0.040)		(0.042)	(0.044)
Responsibility del		-0.027	-0.005		-0.087***	-0.074***
		(0.026)	(0.027)		(0.028)	(0.028)
Attention check		0.169	0.109		0.178	0.095
		(0.256)	(0.315)		(0.317)	(0.339)
Belief FREQ lie			-4.579***			-3.463***
			(0.321)			(0.333)
Constant	-0.773***	-0.979**	1.656***	-1.054***	-0.097	1.997***
	(0.216)	(0.450)	(0.536)	(0.246)	(0.491)	(0.556)
Observations	1,673	1,673	1,673	1,673	1,673	1,673

Notes: Multinominal logit on all three decisions. Robust standard errors in parentheses. The dependent variable takes the value 0 for lying (reference category), 1 for truth-telling, and 2 for delegating. The main independent variable of interest is the position dummy in the game with P1 in D1 as reference category. For gender, female is reference category, and other gender includes "non-binary or other" and "prefer not to answer", with 21 and 10 observations, respectively. Age is measured in years. Risk is measured by the question "Are you generally a person who is willing to take risks, or do you try to avoid taking risks?" on a scale from 0 (lowest willingness to take risks) to 10. Likes to be delegate is measured by the question "Are you generally a person who likes that others delegate decisions to you?" on a scale from 0 (lowest preference to be a delegate) to 10. Attention check is a binary variable that takes the value 0 (1) if the answer to the attention check was correct (wrong). Our measures for need for control, perceptions of responsibility, and belief FREQ lie are described in Appendix B. *, ** and *** document significance at the 10-, 5- and 1-percent level, respectively.

measures, and columns (3) and (6) add the belief about the percentage of subjects that lie in the respective treatment.[16]

Table 5 shows that the dummy for P1 in D2 is always insignificant, both for telling the truth (columns (1) to (3)) and for delegating (columns (4) to (6)). Hence, the

[16]Interactions with gender are never significant and therefore omitted.

behavior of P1s in D2 does not differ significantly from the behavior of P1s in D1. By contrast, while P2s in D2 do not tell the truth more often compared to lying than P1s in D1, they delegate significantly less often ($p = 0.003$ in column (5)). Results are qualitatively the same when comparing P2s in D2 to P1s in D2, i.e., there is no significant difference in truth-telling, but P2s in D2 delegate significantly less often ($p = 0.001$ in the specification mirroring column (5)) compared to P1s in D2.[17]

Compared to the reference category "lying", the controls gender, age, and the attitude towards risk have the same impact on truth-telling and delegating: Gender is insignificant throughout, and both truth-telling and delegating increase with age. The main insights from our personality measures are as follows: First, the willingness to take risk is negative for truth-telling (significant at the 1%-level) and delegating (though only marginally significant at the 10%-level or insignificant), that is, less risk averse people lie more often. People who feel a higher responsibility for the consequences of their behavior lie less often, with significantly positive coefficients at the 1%-level for both truth-telling and delegating. A higher responsibility for the behavior of delegates has no impact on the choice between truth-telling and lying, but reduces the probability of delegating, significant at the 1%-level. Overall, these results show that the proxies for personality used in our survey predict the behavior well in the expected directions.

5.3 Distributions of θ and τ

In this section, we use our experimental data for estimating distributions of lying costs θ and the dilution effect τ. Our estimates are based on Beta distributions, which cover a wide range of different distributions (see the discussion following Proposition 2) and are characterized by only two parameters α and β. Thus, we need to estimate the parameters $\alpha_\theta, \beta_\theta$ of the distribution function F, and α_τ, β_τ of the distribution function G. The estimates are based on the frequencies for the different options as shown in Table 1, aggregated over the three treatments. For instance, the frequency of the last players in the line who lie is 0.551 in treatment ND, 0.575 in treatment D1, and 0.544 in treatment D2. We weigh these probabilities with the respective number of people in

[17] Results available on request.

each treatment. This gives us the aggregated empirical frequencies presented in Table 6.

Table 6: Empirical frequencies and calibrated model probabilities

	Calculation of empirical frequency	Empirical	Model
Last player in line lies	$F(\widetilde{\theta}) = \frac{339 \cdot 0.551 + 344 \cdot 0.575 + 325 \cdot 0.544}{339 + 344 + 325}$	0.557	0.557
Second-to-last player lies	$\mathbb{E}[F(\underline{\theta}_1 \vert \tau_1)] = \frac{244 \cdot 0.408 + 280 \cdot 0.469}{244 + 280}$	0.441	0.429
Second-to-last player delegates	$\mathbb{E}[F(\overline{\theta}_1 \vert \tau_1)] - \mathbb{E}[F(\underline{\theta}_{1,2} \vert \tau_1)] = \frac{148 \cdot 0.248 + 106 \cdot 0.178}{148 + 106}$	0.219	0.241
Third-to-last player lies	$\mathbb{E}[F(\underline{\theta}_{1,3} \vert \tau_1)]$	0.389	0.426
Third-to-last player delegates	$\mathbb{E}[F(\overline{\theta}_1 \vert \tau_1)] - \mathbb{E}[F(\underline{\theta}_{1,3} \vert \tau_1)]$	0.251	0.244

Notes: The table reports empirical lying and delegation frequencies as well as their calibrated model probabilities. For the last player in line, the model's lying probability is assumed to be equal to the empirical frequency.

These frequencies from our expriment and the probabilities from our model are used to estimate the parameters α and β for the two kinds of distributions. We assume that $W = \widetilde{\theta} = 0.557$, and then set β_θ to ensure that the probability of lying for the last player in the line is $F(\widetilde{\theta}) = 0.557$. The remaining parameters $\alpha_\theta, \alpha_\tau, \beta_\tau$ are then determined by minimizing the root mean square of the difference of frequencies and model probabilities, where each of the four remaining differences from Table 6 carries the same weight.

The estimated Beta distributions are shown in Figure 3. The density of θ with calibrated parameters $\hat{\alpha}_\theta = 15.613$ and $\hat{\beta}_\theta = 13.164$ resembles a normal distribution, which is in line with the literature showing a large heterogeneity of lying costs (see Schudy et al. (2023) for an estimation of individual lying costs). By contrast, the density of τ with parameters $\hat{\alpha}_\tau = 40$, $\hat{\beta}_\tau = 2$ is strongly right-skewed and concentrated in an area close to 1, with an expectation of 0.95. While our analysis thereby supports the existence of a dilution of responsibility, it also suggests that the effect is not very pronounced. The previous literature kept silent about the size of the dilution effect, as the probability that selfish allocations are actually implemented was the same with and without delegation (which means that even a very small dilution effect triggers delegation).

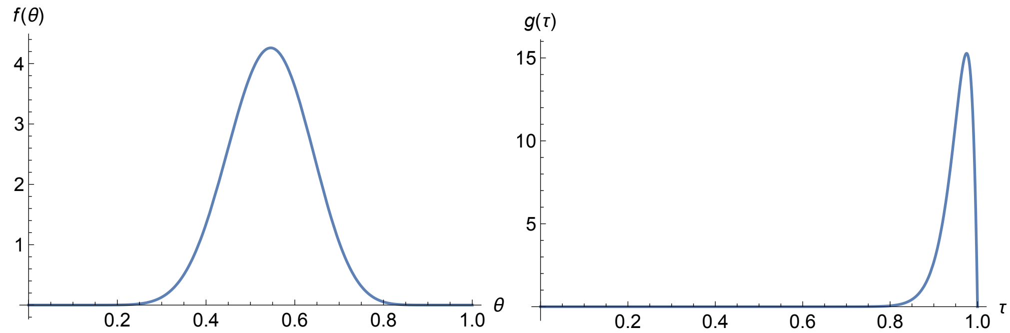

Figure 3: Densities of estimated beta distributions of θ (left) and τ (right). The mean and standard deviation are 0.543 and 0.091 for θ and 0.952 and 0.032 for τ, respectively.

6 Conclusion

Delegating decisions within and across firms is a useful tool for enhancing motivation, increasing the speed of decisions, and exploiting comparative advantages, but may also increase the frequency of morally questionable decisions by reducing moral costs. While the immoral decisions mentioned in the introduction, besides others, are intensively debated in public, we are unaware of theoretical models or experiments that analyze whether delegation opportunities lead to more immoral decisions. Our paper provides a step toward closing this research gap.

In our model, people have private information on their moral costs and their reduction of these costs when decisions are taken by delegates (dilution effect). The overall impact of delegation depends on (i) whether more people switch from moral or immoral decisions to delegating, and (ii) how delegates behave. Our model shows how these two effects interact, and that the result depends on the distributions of types. The pure existence of a dilution effect as already identified by Hamman et al. (2010) is thus no sufficient condition for a detrimental impact of delegation. However, we do find some instructive sufficient conditions for the distribution of moral costs where this holds indeed, regardless of the distribution of the dilution effect. By contrast, a lower probability of immoral decisions can only occur when both distributions fulfil certain conditions.

In our experiment, subjects decide on their potentially inflated reports about the outcome of a lottery. Our parsimonious experimental approach allows us to identify the dilution of responsibility and its consequences. This seems infeasible with field data because, even when delegation is used to facilitate morally questionable decisions, this is not necessarily related to the dilution effect of internal moral costs, but may be motivated by lower detection probabilities or lower legal consequences of detection. In addition, these motives could hardly be disentangled from the bright sides of delegation.

The crucial features of our experiment are that we allow for multi-step delegation, and that those who delegate do not know what their potential delegates will do. We first support previous findings on the presence of the dilution effect, since about 22.5% of our subjects delegate. This could hardly be rationalized without a dilution effect, given that subjects can implement their desired outcome with certainty when deciding on their own. Most importantly, we do not find a significant difference in the frequency of lying with more or less delegation opportunities.

There are two, potentially complementary, channels why delegation may lead to more lying. The first channel is that predominantly those subjects who would have reported honestly delegate. In our experiment, this requires that the ratio of truthful LOST reports over the sum of LOST and WON reports would be lowest for those players who cannot delegate. However, we observe a small though not significant effect in the opposite direction, as the ratios are 44.9% for Players 1 in ND (players who cannot delegate) compared to 45.8% for players 1 in D1, and 48.1% for players 1 in D2 (players who can delegate).

The second potential channel, which is not accounted for in our model, is that delegates lie more often compared to initial decision makers, for instance because they interpret delegation as a request to lie. To check this, we need to compare the behavior of Players 2 in D1, Players 2 in D2, and Players 3 in D2 to the behavior of Players 1 in all three treatments. The ratios of WON reports over the sum of LOST and WON reports is 53.9% for Players 1 (aggregated over all three treatments), compared to 56.3%, (aggregated over Players 2 in D1, Players 2 in D2, and Players 3 in D2). This difference is not significant with $p = 0.176$ (χ^2-tests). As neither of these two channels seems to play a role in our experiment, we do not find significant treatment effects.

We also use our data for estimating distributions over lying costs and the the dilution effect that fit our data best. Restricting attention to Beta distributions, we find a large heterogeneity of lying costs, while the dilution effect is rather small for the majority of players. Previous papers could not provide insights on the size of the dilution effect, as they neglect the downside of delegation from the principal's point of view, which is that delegates might not act in line with the principal's preferences.

Our study is subject to several limitations that offer avenues for future research. First, we have chosen lying on the outcome of a lottery as morally questionable decision. Previous papers on delegation have used variants of the dictator game (Hamman et al., 2010; Bartling and Fischbacher, 2012; Gawn and Innes, 2019b,a). Future research could examine multi-step delegation using allocation decisions to consider the effect of inequity aversion on delegation. Second, given that delegation offers many advantages, one could examine whether there are interaction effects between different motives to delegate. In particular, people who delegate for efficiency reasons may be more or less inclined to also delegate to save on moral costs. Third, our estimates of types are restricted to Beta distributions, and future research could examine lying costs and the dilution effect at the individual level. One step into this direction is taken by Schudy et al. (2023) who examine individual internal lying and reputation costs.

References

Abeler, J., D. Nosenzo, and C. Raymond (2019). Preferences for truth-telling. *Econometrica 87*(4), 1115–1153.

Aghion, P. and J. Tirole (1997). Formal and real authority in organizations. *Journal of Political Economy 105*(1), 1–29.

Akcigit, U., H. Alp, and M. Peters (2021). Lack of selection and limits to delegation: Firm dynamics in developing countries. *American Economic Review 111*(1), 231–275.

Argenton, C., J. Potters, and Y. Yang (2023, sep). Receiving credit: On delegation and responsibility. *European Economic Review 158*, 104522.

Bartling, B. and U. Fischbacher (2012). Shifting the blame: On delegation and responsibility. *The Review of Economic Studies 79*(1), 67–87.

Bicchieri, C., E. Dimant, and S. Sonderegger (2019). It's not a lie if you believe it: On norms, lying, and self-serving belief distortion. Technical report, CeDEx Discussion Paper Series.

Chan, J., M. Selden, and N. Pun (2020). *Dying for an iPhone: Apple, Foxconn, and the Lives of China's Workers*. Haymarket Books.

Charness, G., C. Blanco-Jimenez, L. Ezquerra, and I. Rodriguez-Lara (2019). Cheating, incentives, and money manipulation. *Experimental Economics 22*(1), 155–177.

Chen, D. L., M. Schonger, and C. Wickens (2016). oTree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance 9*, 88–97.

Conrads, J., B. Irlenbusch, R. M. Rilke, and G. Walkowitz (2013). Lying and team incentives. *Journal of Economic Psychology 34*, 1–7.

Dato, S., E. Feess, and P. Nieken (2019). Lying and reciprocity. *Games and Economic Behavior 118*, 193–218.

De Rijk, A. E., P. M. L. Blanc, W. B. Schaufeli, and J. De Jonge (1998). Active coping and need for control as moderators of the job demand-control model: Effects on burnout. *Journal of Occupational and Organizational Psychology 71*(1), 1–18.

Demski, J. S. and D. E. M. Sappington (1987). Delegated expertise. *Journal of Accounting Research 25*(1), 68.

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. G. Wagner (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association 9*(3), 522–550.

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

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

Feess, E. and F. Kerzenmacher (2018). Lying opportunities and incentives to lie: Reference dependence versus reputation. *Games and Economic Behavior 111*, 274–288.

Feess, E., F. Kerzenmacher, and Y. Timofeyev (2022). Utilitarian or deontological models of moral behavior – what predicts morally questionable decisions? *European Economic Review 149*, 104264.

Fehr, E., H. Herz, and T. Wilkening (2013). The lure of authority: Motivation and incentive effects of power. *American Economic Review 103*(4), 1325–1359.

Fischbacher, U. and F. Föllmi-Heusi (2013). Lies in disguise—anexperimental study on cheating. *Journal of the European Economic Association 11*(3), 525–547.

Gawn, G. and R. Innes (2019a). Lying through others: Does delegation promote deception? *Journal of Economic Psychology 71*, 59–73.

Gawn, G. and R. Innes (2019b). Who delegates? Evidence from dictator games. *Economics Letters 181*, 186–189.

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

Gneezy, U., A. Kajackaite, and J. Sobel (2018). Lying aversion and the size of the lie. *American Economic Review 108*(2), 419–453.

Gurdal, M. Y., J. B. Miller, and A. Rustichini (2013). Why blame? *Journal of Political Economy 121*(6), 1205–1247.

Hamman, J. R., G. Loewenstein, and R. A. Weber (2010). Self-interest through delegation: An additional rationale for the principal-agent relationship. *American Economic Review 100*(4), 1826–1846.

Huang, L., J.-S. Song, and R. Swinney (2022). Managing social responsibility in multitier supply chains. *Manufacturing & Service Operations Management 24*(6), 2843–2862.

Innes, R. (2022). Does deception raise or lower lie aversion? experimental evidence. *Journal of Economic Psychology 90*, 102525.

Jacobsen, C., T. R. Fosgaard, and D. Pascual-Ezama (2017). Why do we lie? A practical guide to the dishonety literature. *Journal of Economic Surveys 32*(2), 357–387.

Kajackaite, A. and U. Gneezy (2017, mar). Incentives and cheating. *Games and Economic Behavior 102*, 433–444.

Kocher, M. G., S. Schudy, and L. Spantig (2018). I lie? We lie! Why? Experimental evidence on a dishonesty shift in groups. *Management Science 64*(9), 3995–4008.

Krupka, E. L. and R. A. Weber (2013). Identifying social norms using coordination games: Why does dictator game sharing vary? *Journal of the European Economic Association 11*(3), 495–524.

Mitchell, N. J. (2021). *Why Delegate?* Oxford University Press Inc.

Muehlheusser, G., A. Roider, and N. Wallmeier (2015). Gender differences in honesty: Groups versus individuals. *Economics Letters 128*, 25–29.

Oexl, R. and Z. J. Grossman (2012). Shifting the blame to a powerless intermediary. *Experimental Economics 16*(3), 306–312.

Reuters (2020). Volkswagen says diesel scandal has cost it 31.3 billion euros. www.reuters.com/article/us-volkswagen-results-diesel-idUSKBN2141JB, retrieved on January 13, 2023.

Schudy, S., S. Grundmann, and L. Spantig (2023). Individual preferences for truth-telling. Working paper.

Yukl, G. and P. P. Fu (1999). Determinants of delegation and consultation by managers. *Journal of Organizational Behavior 20*(2), 219–232.

Appendix A: Proofs

Proof of Proposition 2. *Part (i).* With $\widetilde{\theta}$ the median and F symmetric, we have $F(\widetilde{\theta}) = 1 - F(\widetilde{\theta}) = 1/2$. Using Equation (2), it therefore suffices to show that $\mathbb{E}\left[F\left(\overline{\theta}_1|\tau_1\right)\right] - \widetilde{\theta} \geq \widetilde{\theta} - \mathbb{E}\left[F(\underline{\theta}_1|\tau_1)\right]$. Note that

$$\underline{\theta}_1 = \frac{1 - F(\widetilde{\theta})}{1 - \tau_1 F(\widetilde{\theta})} \cdot \widetilde{\theta} = \frac{1/2}{1 - \tau_1/2} \cdot \widetilde{\theta} = \frac{\widetilde{\theta}}{(2 - \tau_1)}$$

$$\overline{\theta}_1 = \frac{\widetilde{\theta}}{\tau_1}.$$

If $\overline{\theta}_1 > 1$, then player 1 will never tell the truth, but prefer delegation to truth-telling with probability 1. In this case, P_2 in (1) simplifies to $P_2 = \mathbb{E}[F(\underline{\theta}_1|\tau_1)] + (1 - \mathbb{E}[F(\underline{\theta}_1|\tau_1)]) \cdot F(\widetilde{\theta}) > F(\widetilde{\theta}) = P_1$.

Consider now the case, where $\overline{\theta}_1 < 1$. Because F is symmetric around $\widetilde{\theta}$, it suffices to show that the length of the interval $[\widetilde{\theta}, \overline{\theta}_1]$ is greater than the length of $[\underline{\theta}_1, \widetilde{\theta}]$, i.e., $\frac{\widetilde{\theta}}{(2-\tau_1)} - \widetilde{\theta} \geq \widetilde{\theta} - \frac{\widetilde{\theta}}{\tau_1}$. Re-arranging and simplifying gives

$$\begin{aligned}\frac{1}{(2-\tau_1)} + \frac{1}{\tau_1} - 2 &= \frac{2 - 2(2-\tau_1)\tau_1}{(2-\tau_1)\tau_1} \\ &= 2 \cdot \left(\frac{1 - (2-\tau_1)\tau_1}{(2-\tau_1)\tau_1}\right) \geq 0,\end{aligned}$$

where the last-step follows from $(2 - \tau_1)\tau_1$ attaining its maximum on $[0, 1]$ at 1. Since this is positive for $\tau_1 \in [0, 1]$, it also positive for all distributions of τ_1.

Part (ii). Using that $F(\underline{\theta}_1|\tau_1) = \frac{1-\widetilde{\theta}}{1-\tau_1\widetilde{\theta}} \cdot \widetilde{\theta}$ and $F(\overline{\theta}_1|\tau_1) = \min\left(\frac{\widetilde{\theta}}{\tau_1}, 1\right) \geq \widetilde{\theta}$ and inserting in (1) gives

$$\begin{aligned}P_2(\tau_1) - P_1 &= \underline{\theta}_1 + (\min(\overline{\theta}_1, 1) - \underline{\theta}_1) \cdot \widetilde{\theta} - \widetilde{\theta} \\ &= \frac{1-\widetilde{\theta}}{1-\tau_1\widetilde{\theta}} \cdot \widetilde{\theta} + \left(\min\left(\frac{\widetilde{\theta}}{\tau_1}, 1\right) - \frac{1-\widetilde{\theta}}{1-\tau_1\widetilde{\theta}} \cdot \widetilde{\theta}\right) \widetilde{\theta} - \widetilde{\theta} \\ &= \frac{(1-\tau_1)^2\widetilde{\theta}}{(1-\tau_1\widetilde{\theta})} \cdot \min\left(\frac{\widetilde{\theta}}{\tau_1}, 1\right) \geq 0.\end{aligned}$$

Since this is positive for all $\tau_1 > 0$, it is also positive for all distributions of τ_1.

Part (iii). Truth-telling is strictly dominated by delegation as $\tau_1 \to 0$, so that $\overline{\theta}_1 \to 1$. Formally, $\lim_{\tau_1\to 0} F(\overline{\theta}_1|\tau_1) = \lim_{\tau_1\to 0} F\left(\frac{\widetilde{\theta}}{\tau_1}\right) = 1$. Likewise, $\lim_{\tau_1\to 0} F(\underline{\theta}_1|\tau_1)$ $= \lim_{\tau_1\to 0} F\left(\frac{1-P_1}{1-\tau_1 P_1}\,\widetilde{\theta}\right) = F((1-F(\widetilde{\theta}))\widetilde{\theta})$. This gives (cf. (2))

$$\begin{aligned}\lim_{\tau_1\to 0} P_2(\tau_1) - P_1 &= (1-F(\widetilde{\theta}))F(\widetilde{\theta}) - (F(\widetilde{\theta}) - F((1-F(\widetilde{\theta}))\widetilde{\theta}))(1-F(\widetilde{\theta}))\\ &= (1-F(\widetilde{\theta}))F((1-F(\widetilde{\theta}))\widetilde{\theta}) > 0,\end{aligned}$$

where the strict inequality follows from $\widetilde{\theta} \in (0,1)$ and F being defined on $[0,1]$. Because of the continuity of $P_2(\tau_1)$ in τ_1, there exists $\widetilde{\tau}_1$, dependent on the distribution of θ and on $\widetilde{\theta}$, such that $P_2(\tau_1) - P_1 > 0$, for all $\tau_1 \le \widetilde{\tau}_1$. Hence, if $\tau_1 \le \widetilde{\tau}_1$ with probability 1, then $P_2 - P_1 > 0$.

Part (iv). For any sequence of $\tau_1^{(n)}$ with $\lim_{n\to\infty} \mathbb{E}[\tau_1^{(n)}] = 0$, it follows from the convexity of $\widetilde{\theta}/\tau_1^{(n)}$ and Jensen's inequality that $\mathbb{E}[\widetilde{\theta}/\tau_1^{(n)}] \ge \widetilde{\theta}/\mathbb{E}[\tau_1^{(n)}] \to \infty$ as $n \to \infty$. As a consequence, $\lim_{n\to\infty} \mathbb{E}[F(\overline{\theta}_1|\tau_1^{(n)})] = 1$ and $\lim_{n\to\infty} \mathbb{E}[F(\underline{\theta}_1|\tau_1^{(n)})] = F((1-P_1)\widetilde{\theta})$, so that $\lim_{n\to\infty} P_2^{(n)} = F((1-P_1)\widetilde{\theta}) + (1-F((1-P_1)\widetilde{\theta}))F(\widetilde{\theta}) > F(\widetilde{\theta}) = P_1$. On the other hand, for any sequence of $\tau_1^{(n)}$ with $\lim_{n\to\infty} \mathbb{E}[\tau_1^{(n)}] = 1$, it follows that $P_2^{(n)} \to P_1$, as $n \to \infty$. If $P_2(\tau_1) - P_1$ is monotone in τ_1, then it is positive and the claim follows. If $P_2(\tau_1) - P_1$ is convex in τ_1, but not positive for all $\tau_1 \in [0,1]$, then there exists τ^*, such that $P_2(\tau^*) = P_1(\tau^*)$ and $P_2(\tau_1) - P_1(\tau_1) > 0$ for all $\tau_1 < \tau^*$. The claim then follows from the convexity of $P_2(\tau_1) - P_1(\tau_1)$ and Jensen's inequality. ■

Proof of Lemma 1. Recall that $\underline{\theta}_{j,n}(\tau_j) = W\dfrac{1-P_{n-j}}{1-\tau_j P_{n-j}}$, where P_{n-j} is the probability of reporting W conditional on player j delegating. First, observe that $\underline{\theta}_{j,n}(\tau_j)$ is monotone decreasing in P_{n-j} by checking that

$$\frac{\partial}{\partial q}\frac{1-q}{1-\tau_j q} < 0. \tag{3}$$

When moving from an n-player game to an $n+1$-player game, the only change in $\underline{\theta}_{j,n}(\tau_j)$ is the change in probability from P_{n-j} to P_{n+1-j}. If this probability increases, then according to (3), $\underline{\theta}_{j,n}$ decreases in n. Likewise, if the probability decreases. ■

Proof of Proposition 3. We first prove part *(iii)* and then show how parts *(i)* and *(ii)* follow from *(iii)*.

Part (iii) Without loss of generality we consider player 1. The proof proceeds by induction. Player 1 delegates in an n-player game if

$$\underline{\theta}_{1,n} = W \cdot \frac{1 - P_{n-1}}{1 - \tau_1 P_{n-1}} < \theta_1 \leq \frac{W}{\tau_1} = \overline{\theta}_1.$$

If $P_n > P_{n-1}$ (the induction argument), then by Lemma 1 the lower bound for delegating decreases, i.e., $\underline{\theta}_{1,n} > \underline{\theta}_{1,n+1}$. Recall that the probability of subject 1 reporting L is determined from $W - \tau_1\theta_1 < 0$, which does not depend on the number of players in the game. In other words, the sum of the probabilities of reporting W or delegating does not depend on n. Write this as

$$\mathbb{P}_{n+1}(W_1|\tau_1) + \mathbb{P}_{n+1}(D_1|\tau_1) = \mathbb{P}_n(W_1|\tau_1) + \mathbb{P}_n(D_1|\tau_1) = F(\overline{\theta}_1|\tau_1) = F\left(\frac{W}{\tau_1}\right), \quad (4)$$

where W_1 and D_1 denote the events that player 1 reports W or delegates, respectively, and $\mathbb{P}_n$ denotes the probability measure in a game with n players. Also recall that

$$P_n = \mathbb{P}_n(W_1) + \mathbb{P}_n(D_1) \cdot P_{n-1} \quad (5)$$

and that

$$\mathbb{P}_{n+1}(W_1|\tau_1) = F(\underline{\theta}_{1,n}|\tau_1) = F\left(W \cdot \frac{1 - P_n}{1 - \tau_1 P_n}\right). \quad (6)$$

The difference in the probability of lying is given as

$$\begin{aligned}
P_{n+1} - P_n &\overset{(5)}{=} \mathbb{P}_{n+1}(W_1) + \mathbb{P}_{n+1}(D_1)P_n - (\mathbb{P}_n(W_1) + \mathbb{P}_n(D_1)P_{n-1}) \\
&= (\mathbb{P}_{n+1}(W_1) + \mathbb{P}_{n+1}(D_1))P_n - (\mathbb{P}_n(W_1) + \mathbb{P}_n(D_1))P_{n-1} \\
&\quad + \mathbb{P}_{n+1}(W_1)(1 - P_n) - \mathbb{P}_n(W_1)(1 - P_{n-1}) \\
&\overset{(4),\,(6)}{=} \mathbb{E}\left[F\left(\frac{W}{\tau_1}\right)\right] P_n - \mathbb{E}\left[F\left(\frac{W}{\tau_1}\right)\right] P_{n-1} \\
&\quad + \mathbb{E}\left[F\left(W\frac{1 - P_n}{1 - \tau_1 P_n}\right)\right](1 - P_n) - \mathbb{E}\left[F\left(W\frac{1 - P_{n-1}}{1 - \tau_1 P_{n-1}}\right)\right](1 - P_{n-1}).
\end{aligned}$$

Since $P_n > P_{n-1}$ by induction, the difference can be written as

$$P_{n+1} - P_n = \int_{P_{n-1}}^{P_n} \frac{\partial}{\partial q} \mathbb{E}\left[\left(F\left(\frac{W}{\tau_1}\right) - F\left(W\frac{1-q}{1-\tau_1 q}\right)\right) q + F\left(W\frac{1-q}{1-\tau_1 q}\right)\right] dq,$$

and it remains to show that the integrand is positive. Because of the induction structure, we have that there exists some $q_0 \leq P_{n-1}$, for which the integrand is positive, therefore it is sufficient to show that the integrand is increasing in q.

The integrand is given as

$$\begin{aligned}
&\frac{\partial}{\partial q} \mathbb{E}\left[\left(F\left(\frac{W}{\tau_1}\right) - F\left(W\frac{1-q}{1-\tau_1 q}\right)\right) q + F\left(W\frac{1-q}{1-\tau_1 q}\right)\right] \\
&\qquad = \mathbb{E}\left[F\left(\frac{W}{\tau_1}\right) - F\left(W\frac{1-q}{1-\tau_1 q}\right) - \frac{1-\tau_1}{1-\tau_1 q} W \frac{1-q}{1-\tau_1 q} F'\left(W\frac{1-q}{1-\tau_1 q}\right)\right].
\end{aligned}$$

The first part $\mathbb{E}\left[F\left(\frac{W}{\tau_1}\right) - F\left(W\frac{1-q}{1-\tau_1 q}\right)\right]$ increases at a rate $\mathbb{E}\left[F'\left(W\frac{1-q}{1-\tau_1 q}\right)\frac{W(1-\tau_1)}{(1-\tau_1 q)^2}\right]$.

For the second part, we have

$$\begin{aligned}
&\frac{\partial}{\partial q} \mathbb{E}\left[\frac{W(1-\tau_1)(1-q)}{(1-\tau_1 q)^2} F'\left(W\frac{1-q}{1-\tau_1 q}\right)\right] \\
&= \mathbb{E}\left[W(1-\tau_1)\frac{2\tau_1 - 1 - \tau_1 q}{(1-\tau_1 q)^3} F'\left(W\frac{1-q}{1-\tau_1 q}\right) - \frac{W(1-\tau_1)^2(1-q)}{(1-\tau_1 q)^4} F''\left(W\frac{1-q}{1-\tau_1 q}\right)\right].
\end{aligned}$$

As a consequence, using that $\mathbb{E}\left[F''\left(W\frac{1-q}{1-\tau_1 q}\right)\right] \geq 0$,

$$\begin{aligned}
&\frac{\partial}{\partial q} \mathbb{E}\left[F\left(\frac{W}{\tau_1}\right) - F\left(W\frac{1-q}{1-\tau_1 q}\right) - \frac{1-\tau_1}{1-\tau_1 q} W \frac{1-q}{1-\tau_1 q} F'\left(W\frac{1-q}{1-\tau_1 q}\right)\right] \\
&\quad = \mathbb{E}\left[F'\left(W\frac{1-q}{1-\tau_1 q}\right)\frac{W(1-\tau_1)}{(1-\tau_1 q)^2}\left[1 - \frac{2\tau_1 - 1 - \tau_1 q}{1-\tau_1 q} + \frac{(1-\tau_1)(1-q)}{(1-\tau_1 q)^2}\frac{F''}{F'}\right]\right] \\
&\quad = \mathbb{E}\left[F'\left(W\frac{1-q}{1-\tau_1 q}\right)\frac{W(1-\tau_1)}{(1-\tau_1 q)^2}\left[\frac{2(1-\tau_1)}{1-\tau_1 q} + \frac{(1-\tau_1)(1-q)}{(1-\tau_1 q)^2}\frac{F''}{F'}\right]\right] \\
&\quad \geq 0.
\end{aligned}$$

Part *(i)*. That $P_2 > P_1$ follows from Proposition 2. Because F is symmetric and unimodal, it has its unique maximum at the median W. As a consequence, the density of F is non-decreasing for all values below W, i.e. $F''(\hat{\theta}) \geq 0$ for $\hat{\theta} \leq W$.

Part *(ii)*. This follows directly from *(iii)*. ■

Appendix B: Construction of control variables

The following describes the construction of the control variables that are built from more than one item.

Belief FREQ lie is built from beliefs about the percentages of other subjects who report W, L, and delegate (if possible) in their respective roles when the lottery was lost. In treatment *ND*, *Belief FREQ lie* captures the beliefs about the share of subjects reporting W. In treatment *D1*, *Belief FREQ lie* captures the beliefs about the share of P1 reporting W plus beliefs about the share of P1 who delegate multiplied by beliefs about the share of P2 reporting W. In treatment *D2*, *Belief FREQ lie* is calculated as in treatment *D1*, plus beliefs about the share of P2 who delegate multiplied by beliefs about the share of P3 reporting W.

Need for control is constructed from the following four items (De Rijk et al. 1998) measured on a scale from 0 (lowest preference to be in control) to 10. The four items are "I prefer giving orders instead of receiving them.", "I prefer having control over what I do and the way I do it.", "I prefer doing my own planning.", and "I prefer being able to set the pace of my tasks.". *Need for control* gives the average of the four items.

Responsibility own is constructed from the following two items measured on a scale from 0 (lowest responsibility) to 10. The two items are "If I make a decision and lie to another person, I am responsible for this lie." and "If I make a decision that harms a third person, I am responsible for this harm.". *Responsibility own* gives the average of the two items.

Responsibility del is constructed from the following two items measured on a scale from 0 (lowest responsibility) to 10. The two items are "If I delegate a decision and the delegate lies to another person, I am responsible for this lie." and "If I delegate a decision and the delegate's decision harms a third person, I am responsible for this harm.". *Responsibility del* gives the average of the two items.