Lying Through Others:

Does Delegation Promote Deception?

Glynis Gawn Robert Innes *

April 2018

Abstract

How do agency relationships affect an individual's willingness to lie for monetary advantage? Does lie aversion decline if a lie (or truth) is sent through an agent, rather than sent directly by the individual? In a recent paper, Erat (2013) shows that a significant proportion of his subjects prefer to delegate a deception decision. We present experiments designed to focus on one of several possible explanations for this intriguing behavior – that delegation reduces lie aversion. The experiments reveal that subjects are more willing to lie through a delegate than to lie directly despite controlling for potential effects of delegated decision-making on preferences over payoffs, probabilities of actions, and/or the desire to avoid taking a decision.

Keywords: Delegation, Deception, Lie Aversion

JEL Classifications: D03

*Economics, University of California, Merced. School of Social Sciences, Humanities & Arts, U.C. Merced, 5200 N. Lake Rd., Merced, CA 95343. Emails: ggawn@ucmerced.edu, rinnes@ucmerced.edu. Phone: (209) 228-7742.

1 Introduction

Agency relationships are ubiquitous in economic interactions. Outsourced suppliers and subcontractors produce for contracting companies. Employees act for firms. Intermediaries act on behalf of clients in employment and professional services. In many of these settings, agents make decisions with ethical overtones that affect economic costs and benefits to their principals. Outsourced suppliers may misrepresent questionable labor practices (e.g., Nike in the 1990's) or noxious product contents (e.g., Lumber Liquidators in 2015). External auditors may misrepresent financial conditions (e.g., Enron, Worldcom). Employment agencies may hire illegal immigrants as house help for clients (e.g., Mitt Romney, as discussed in Erat, 2013).[1] One type of moral decision – relevant to many of these examples – involves deception or lying for monetary advantage.

In a recent paper, Erat (2013) is the first to study delegation in a deception game. In his experiment, each Sender privately observes the outcome of a die roll and transmits a message indicating the outcome to an uninformed Receiver. Based only on the message and no knowledge of payoffs, the Receiver then reports a number he or she thinks was rolled. The Sender obtains a higher payoff if the Receiver reports an incorrect roll, and the Receiver obtains a higher payoff if the true roll is reported. Anticipating that the Receiver will report according to the only available cue – the message – the Sender thus benefits materially from a lie and the Receiver is harmed. Erat's (2013) innovation is to give Senders the option either to delegate the message decision to a randomly selected player, or to make the decision themselves. If opting to delegate, the message chosen by the randomly selected player is sent to the Sender's own Receiver. In this experiment, Erat (2013) obtains a fascinating result: a significant fraction of subjects opt to delegate the deception decision, rather than choose

[1] See, for example, O'Rourke (1997) on Nike, McLean & Elkind (2003) on Enron, 60 Minutes (2015) on Lumber Liquidators, and Friedman (2011) on Romney.

a message themselves.

There are a few possible explanations for this behavior. Delegation may offer subjects a cognitive opt-out, enabling them to avoid a difficult decision (see Dana et al., 2006, p. 197). Alternately, delegation might attenuate "other regarding preferences" – the preference for the generous payoff option that a truth promotes. If a lie benefits the liar at the expense of the recipient – a "black lie" (Erat & Gneezy, 2012) as in Erat's (2013) experiment – then lying can be motivated by a preference for the self-benefiting payoffs that a lie produces (Hurkens & Kartik, 2009). In dictator games where a dictator chooses an allocation of money between herself and another player, Hamman, Lowenstein & Weber (2010) find evidence that delegation reduces generosity (see also Fershtman & Gneezy, 2001; Drugov, Hamman & Serra, 2014). Principals / dictators choose agents who tend to be more "selfish" than the principals would be when directly choosing payments. [2] In Erat's (2013) experiment, delegation might be favored because it lessens the preference for costly generosity.

A different explanation is stressed by Erat (2013) and is our focus in this paper: If a principal's decision is delegated – made through an agent – then the principal's "pure" aversion to a lie is reduced. In this paper, we present experiments designed to determine whether this effect is at work when controlling for the other possible economic channels by which a delegation option could prove attractive. If you are choosing to lie or not via an agent, such as an outsourced supplier, is your aversion to the lie less than if you were directly

[2]To our knowledge, Fershtman & Gneezy (FG, 2001) are the first to identify a strategic benefit of delegation in reducing the preference for generosity. In an ultimatum game where a Receiver can reject a dictator's allocation offer, FG find that offers are less generous, and the rate at which Receivers reject stingy offers is reduced, when the offer is made via a delegate who is also harmed by a Receiver's rejection. See also Blount (1995). Hamman, Lowenstein & Weber (HLW, 2010) show that delegation leads to allocations that are more favorable to the principal and less favorable to the matched player in dictator games with interactive principal-agent play. Drugov, Hamman & Serra (2014) find that delegation makes principals more selfish in a corruption game where the principal/briber and a bribee can increase their own payoffs at the expense of a third party. Bartling & Fischbacher (2012) show that delegation shifts the blame for selfish decisions. Overall, this literature indicates that delegated (vs. direct) decision making can be advantageous to principals by reducing the demand for "fair" allocations that share money with others.

deciding whether or not to lie? To answer this question, we compare deception decisions made directly by a subject (by choice of a lie vs. a truth) to those made through an agent (by choice of a lying agent vs. a truth-telling agent). In the classroom experiments, we find that delegation increases the propensity for deception, supporting Erat's (2013) interpretation of his findings.

The experiments control for three effects that delegation can have on decisions in a deception game such as Erat's (2013), separate from impacts on lie aversion. First, the experiments preclude decision opt-outs – delegating in order to avoid cognitive and emotional costs of taking a decision – by forcing subjects to make equivalent choices with and without delegation.

Second, using parallel dictator games in our "black lie" experiment, and a "white lie" design in which both players benefit from a lie (Erat & Gneezy, 2012), the experiments control for potential effects of delegation on preferences over allocations.

Third, delegation can bring about two inter-related changes to the decision environment, reducing (1) control over a decision and outcomes and (2) attribution of responsibility for a decision. Often the two are considered one and the same. Philosophers argue that individuals are responsible for outcomes only if they can control them (Nelkin, 2004; Gurdal, Miller & Rustichini, 2013). Ceding of control is a central feature of delegation in the recent experimental literature studying its effects [3] and is characteristic of many delegation situations in practice. However, in many of the examples cited at the outset, delegation need not cede control per se: agents can be selected because they will make decisions that the principals would themselves make if they could immunize themselves from responsibility. [4]

[3] A noteworthy exception is Drugov et al. (2014), who study effects of a powerless intermediary in a corruption game.

[4] A loss of control is also symptomatic of situations posed in the literature on "moral wiggle room" (Dana, Weber & Kuang, 2007). When negative outcomes for a matched player can be due to either nature or a

Even in these cases, delegation may reduce both the extent to which others assign blame to the principal (as in Bartling and Fischbacher, 2012, Coffman, 2011, and Oexl & Grossman, 2013) [5] and a principal's own sense of responsibility for a morally noxious decision. The latter is our focus: How does delegation affect lie aversion when there is no loss of control over decisions and no scope for others to judge or punish? Our experiments avoid a loss of control under delegated vs. direct decision-making by equating probabilities that decisions translate into actions in the two cases.

After a brief overview of our approach and hypotheses in Section 2, the two experiments are presented in Sections 3 and 4 below. Section 5 presents a follow-on experiment that examines robustness of our results to several design features. Section 6 concludes. We note at the outset that we use the term *delegation* to refer to decision-making through another person, even though there is no loss of control; such situations might alternately be called *intermediation*.

2 Overview of Experiments and Hypotheses

How does the process of delegation affect an individual's aversion to lies? We base our experiments on deception interactions in which one person (the Sender) is lying or telling the truth to another person (the Receiver), as in Gneezy's (2005) initial design, rather than

dictator's decision – with an attendant reduction in the probability that the dictator's decision is implemented – then dictators tend to "hide behind nature" and act more selfishly (Dana et al., 2007; Andreoni & Bernheim, 2009). See also Bartling & Fischbacher (2012), where punishment is possible; Bartling & Fischbacher (2012) explain their results using a new measure of responsibility driven by the extent to which a principal is perceived to affect the outcome. An interesting paper by Haisley & Weber (2010) shows that subjects may "hide behind ambiguity"; adding ambiguity to probability distributions of outcomes – even when ultimate probabilities are the same – leads to more selfish decisions. Ambiguity does not reduce control objectively, but may be perceived as doing so, as indicated by the authors. Decisions in groups, or based on team incentives, can make subjects more selfish (Falk & Szech, 2013) and more willing to lie (Conrads et al., 2013). Studies on responsibility alleviation (Charness, 2000) and hidden costs of control (Falk & Kosfeld, 2006) show how reducing an agent's control can elevate moral hazard.

[5] See also Eisenkopf & Fischbacher (2015) who find that Returners in a trust game differentially reward the person who makes the trust decision, whether a delegate or a Sender.

recent variations in which there is no one at the receiving end of the lie/truth (e.g., Gibson, et al., 2013; Fischbacher & Heusi, 2013; Abeler et al., 2014). We adopt a two-sided deception frame because we believe delegation is only likely to be meaningful if there is interpersonal content to decisions.

In each of our experiments, there are two possible payoff allocations, one that is better for the Sender (call it Option 1) and the other that is worse for the Sender (call it Option 2). If a lie is sent on behalf of the Sender, Option 1 is more likely to be implemented; conversely, if a truthful message is sent, Option 2 is more likely. In Experiment 1, the first (Option 1) allocation is worse for the Receiver, so that a lie benefits the Sender at the expense of the Receiver. In Experiment 2, Option 1 is better for both Sender and Receiver. Erat & Gneezy (2012) describe the former situation as a "black lie" and the latter as a "white lie", nomenclature that we will borrow.

The black lie Experiment 1 is a two-by-two between-subject design with two games (deception and dictator) and two choice environments (direct and delegated). Each subject participates in one of the four treatments, and the payoff options are identical across all four. In the direct treatments, Senders choose directly whether to lie or not (deception) and which payoff option to select (dictator). In the delegation treatments, Senders each choose one of two agents (Agent 1 vs. Agent 2) and the decision of that agent is implemented. In the deception game, Agent 1 has lied and is more likely to produce a lie on behalf of the Sender; Agent 2 has told the truth and is more likely to produce a truth on behalf of the Sender. Similarly, in the dictator game, Agent 1 is more likely to produce the Sender-benefiting payoff option and Agent 2 is more likely to produce the generous (Receiver-benefiting) payoff option.

In both treatments, Senders express a preference for a lie vs. a truth (in the deception game) or the selfish vs. generous option (in the dictator game), either by their direct choice

(in the direct treatment) or by their choice of agent (in the delegation treatment). In each case, a truthful/generous (vs. untruthful/selfish) decision produces exactly the same probabilities of ultimate outcomes. This is done by matching probabilities across treatments in order to avoid any differential uncertainty in consequences. For example, in the delegation treatments, the actual decisions of Agents 1 and 2 are known to the Sender; if a Sender chooses the lying Agent 1, the lie is sent with exactly the same probability as when a lie is chosen in the direct treatment. As a result, delegation does not cede any control over decisions.

The prediction of interest is that delegation – implementing a lie through another person rather than directly – reduces lie aversion, attenuating moral preferences by adding distance between the Sender and Receiver. This prediction has three symptoms in our experiments:

> *Hypothesis 1 (the difference test).* A larger fraction of subjects will lie when making a delegated decision (through an Agent) than when making the decision directly.

In principle, evidence for Hypothesis 1 from our black lie Experiment 1 could be attributable to impacts of delegation on preferences over allocations (the effect identified in Hamman et al., 2010). More subjects may lie under delegation because more subjects prefer the "selfish allocation". To net out the latter effects and thereby enable inferences about lie aversion, we have:

> *Hypothesis 2 (the difference-in-difference test).* Delegation will increase the fraction of subjects who lie (by selecting the lying Agent) by more than it increases the fraction of subjects who choose the selfish payoff allocation (by selecting the "selfish" Agent).

Hypothesis 2 follows Gneezy's (2005) approach to identifying treatment effects on lie aversion.

A difference with Gneezy (2005) is that our treatments do not vary payoffs. [6]

Hurkens & Kartik (2009) propose a different approach to disentangling treatment effects on payoff preferences. Because only those who prefer the "selfish" payoff would lie – and these subjects would only tell the truth due to an aversion to lying – Hurkens & Kartik (2009) argue that lie aversion can be measured by conditioning on a selfish choice in the dictator game. Conditional on preferring the self-benefiting payoff allocation, to what extent do subjects lie versus tell the truth? If the conditional proportion of lies falls in one setting versus another, then one can conclude that lie aversion has risen on average.

> *Hypothesis 3 (the Hurkens and Kartik test).* Delegation increases the fraction of selfish subjects who lie.

3 The Black Lie Experiment 1

In Experiment 1, each subject participates in one of two games – deception and parallel dictator – and one of two choice environments, direct and delegation. In each of the four cases, Senders interact with matched Receivers in different classrooms. Our interest is in the Senders. Our baseline deception game is adapted from the Erat & Gneezy (2012) "white lie" experiments, drawing on the Lopez-Perez & Spiegelman (2013) dot framework.

Direct Deception Game. Each Sender observes a dot that is either blue or green. The matched Receiver does not observe the dot. A message is sent to the Receiver on behalf of the Sender indicating the true color of the dot. Two possible messages can be sent:

> **Message GREEN**: I solemnly swear that the dot is GREEN.
>
> **Message BLUE**: I solemnly swear that the dot is BLUE.

[6] The on-line Appendix presents a simple conceptual model to frame the experiments. Allowing for a general distribution of preferences over payoffs (to Sender and Receiver) and lies (vs. truths), the model produces Hypothesis 2 as a test of treatment effects on lie aversion when the Hamman et al. (2010) effect is present.

We specify "strong" messages in order to add salience to the truthfulness of the Sender's choice. Based only on the message he or she receives, the Receiver must Report whether the dot is blue or green. Depending upon the Report, payments are made to the Sender and Receiver. The payments are as follows:[7]

> If the Receiver reports *the true color* (say blue), then payments are \$5 each to Sender and Receiver.
>
> If the Receiver reports *the incorrect color* (say green), then payments are \$7 to the Sender and \$3 to the Receiver.

The Sender earns more with an incorrect Receiver report, giving him/her a monetary incentive to lie by sending an untruthful message (Message Green in the case of a blue dot).

Following the literature (e.g., Gneezy, 2005; Sutter, 2009; Erat, 2013), the Receiver is told that the Sender observes the payments associated with the Report, and nothing more about payments. To avoid strategic considerations, Senders (but not Receivers) are told that "in similar experiments, Receivers almost always report according to the Message they receive". This statement is based on the Erat (2013) experiments where 15 of 16 Receivers report according to the message.[8] In our modified experiment, Receivers report according to the message in 80.4 percent of cases. 74.4 percent of Senders predict that their Receiver will follow their message.[9]

[7]There is no obvious choice of payment options. Our payments are similar to those in Innes & Mitra (2013), where the gain to Sender dishonesty is \$2 and the corresponding Receiver loss is \$3. In our case, the Sender gain and Receiver loss are both \$2, which might be expected to favor more dishonesty (Gneezy, 2005). However, our messages are stronger and the truthful choice is associated with equal payoffs, both of which might be expected to favor more honesty.

[8]In Erat (2013) and Erat & Gneezy (2012), the message and report reflect the outcome of a die roll, namely, a number from one to six. The expanded message space might be expected to increase Receivers' propensity to report according to the messages received, relative to our (two option) space; however, our stronger messages might also be expected to increase Receiver acceptance rates.

[9]At the end of the experiment (after all decisions have been made), predictions are elicited from most Senders in the deception experiment (132 of 181 participants). Each Sender is asked to predict whether

Senders participate in two "situations", one of which is randomly selected for payment. The first situation – Situation K – is implemented with (2/3) probability and the second – Situation L – is implemented with (1/3) probability. In Situation K, the Sender chooses which Message to send, Message Green or Message Blue. In Situation L, the truthful Message (Message Blue, for example) is automatically sent on the Sender's behalf. This structure mimics the probabilities in the delegation treatment, as will become clear in a moment. It implies that a truthful decision in Situation K (Message Blue) prompts a truthful (blue) message with probability one; similarly, an untruthful decision (Message Green) leads to an untruthful message with (2/3) probability and a truthful message with (1/3) probability.

Delegated Deception Game. In the delegation treatment, Senders do not themselves choose the message. Instead, they choose between two Agents (live subjects in a different class). Having observed the same color dot as the Sender, one of the Agents has chosen a Blue Message (truthful, say) and the other has chosen a Green Message (untruthful). The Sender chooses whether to delegate to the "blue" Agent or the "green" Agent (identified as Agent 1 and Agent 2 in the instructions, with labels randomly varied). Once an Agent is selected, the corresponding message choice of the Agent is sent to the Receiver with (2/3) probability and the truthful (blue) message is automatically sent with (1/3) probability.[10] Once Messages are sent, Senders and Receivers are paid according to the same rules as in the direct deception game. We adhere to a probabilistic effect of Agent selection on message

their own Receiver will follow their message (or not) and is paid $1 if the prediction is correct. The Sender predictions and Receiver decisions are similar to those found in the literature. For example, 78 percent of Gneezy's (2005) Receivers follow their message, as do 72 percent in Sutter (2009) and 73 percent in Innes & Mitra (2013). 72.3 percent of Sutter's (2009) Senders predict that their Receivers will follow, as do 73.4 percent in Innes & Mitra (2013).

[10]The Receiver instructions are slightly different in the delegated deception game from those in the direct deception game. As the Message is not chosen directly by the Sender, the instructions indicate that the Message is chosen by another student who observes exactly the same color dot as the Sender and chooses a message to send on behalf of the Sender, knowing the payment options. With the "direct" instructions, 76.4 percent of Receivers follow the Message; in the "delegated" instructions, 84.2 percent follow. The difference is not statistically significant (z=1.047).

choice in order to mimic the uncertainty that delegation yields in practice, but without any differential loss of control across the treatments. Figure 1 summarizes the design.

Agents each choose one Message after being shown a dot and told the complete structure of the game. For example, if shown a blue dot, the Agent is either a "blue Agent" (if she chooses Message Blue) or a "green Agent" (if she chooses Message Green) for Senders who are also shown a blue dot. Agents are not paid based on payoffs in the Sender-Receiver game. Rather they are paid according to the number of Senders who choose them as their Agent, based on their choice of Message. If multiple Agents choose the same Message (say green when the dot is blue), each has an equal probability of being designated as the corresponding (green) Agent in any (blue-dot) Sender questionnaire. Senders are told simply that "the Agents are not paid based on the payments attached to the Messages, but rather according to how many Senders (like you) choose them as their own Agent."[11]

Direct and Delegated Dictator Games. Parallel direct and delegated dictator games are also implemented, each with different subjects. Payoff options are identical to those described above: \$5 each (Option A) vs. \$7 for the Sender and \$3 for the Receiver (Option B). For the delegated dictator game, the Sender/dictator chooses between two Agents, one of whom chose the equal-split Option A (Agent A) and the other of whom chose the unequal Option B (Agent B). In both games, decisions are implemented probabilistically. With (2/3) probability, the Sender chooses the payment option (in the direct treatment) or the Agent (in the delegated treatment). With (1/3) probability, the equal-split (\$5 each) option is automatically chosen. With this design, if a Sender chooses the generous Option/Agent A, then Option A is selected with probability one. If a Sender chooses the "selfish" Option/Agent B,

[11]Delegating dictators are given a similar statement: "The Agents are not paid based on the payment options, but rather according to how many Senders (like you) choose them as their own Agent." In our Agent session for the deception game, participants produced all four possible combinations of dot colors and Message choices (blue-blue, blue-green, green-blue and green-green).

then Option B ($7/$3 to Sender/Receiver) is selected with (2/3) probability and the other ($5 each) Option A is selected with (1/3) probability. In order to match Receiver acceptance decisions in the deception game (following Gneezy, 2005), the "selected" option is implemented with 80 percent probability and the other option is implemented with 20 percent probability. The 80 percent probability corresponds exactly with actual Receiver following decisions. Figure 2 depicts the two dictator treatments. [12]

Matching Probabilities. In the experiment, a subject's decision to lie produces the same probability of an untruthful Message, whether the decision is made directly or by choice of a lying Agent (Figure 1). A subject's decision to be "selfish" in a dictator game produces the same probability of a "selfish" allocation, whether the decision is made directly or by choice of a "selfish" Agent (Figure 2). With an 80 percent probability that Receivers follow messages – the true probability in the experiment – a subject's decision to lie in the deception game also produces the same probability of a "selfish" allocation as does a "selfish" decision in the dictator game. Hence, there are no differences in how choices affect monetary consequences across the treatments.

Logistics. We conduct the Sender side of Experiment 1 in seven lower- and upper-division undergraduate economics classes at U.C. Merced. In each classroom, questionnaires corresponding to the four different treatments are equally mixed and randomly distributed. We randomly vary dot colors (in the deception games), Option labels (in the dictator games) and Agent labels (in the delegation treatments). In total, there are 356 Sender-Dictator/Receiver pairs, with 87 Senders in the direct deception game, 94 in delegated deception, 88 in direct dictator, and 87 in delegated dictator. Receivers are in different classes than paired Senders.

[12] As in the deception games, Receivers in the dictator games are not told anything about dollar payoffs. They are only told that another student has made a decision and, based upon that decision, payments will be made to the two of them.

Agents are from a small Political Science graduate class.

Participation in the experiment is purely voluntary and has no bearing on course assessment. Class rosters are used to avoid a student participating more than once. Subjects are instructed not to communicate with each other and are carefully monitored to ensure privacy. Participant decisions are completely anonymous.[13] Payments are based on registration numbers attached to the questionnaires and are collected by the students one week after the experiment.[14] To avoid any potential subject concern for dollar costs to the experimenters or a desire to "please" (see Lopez-Perez & Spiegelman, 2013; Levitt & List, 2007), we also convey verbally to the students that we are interested in their decisions for our research and that they should make the decisions that they would like to make under the indicated circumstances.

In the Sender sample, the proportion of male subjects exhibits only minor variation across the treatments, ranging from a low of 51.5% (delegated deception) to a high of 57.9% (direct dictator); 52.9% of subjects are male in the other two treatments. By design, the treatments are roughly equally represented in all Sender classrooms.

Results from Experiment 1. Figure 3 and Table 1 present results from Experiment 1, indicating overall percentages of untruthfulness in the deception games (direct and delegated)

[13]The students are told (both verbally and in the written instructions): "The decisions that you make in this experiment will be completely anonymous," with the verbal addition: "No one, including us, will ever know that you are the person making the decisions you make." Decisions are made inside the questionnaire and are not visible to the experimenter when the questionnaire is returned; subjects return their questionnaires to a box at the front of the room when they are finished, with their decision sheets covered over by a front sheet. Participants are seated facing the front of the room with ample space between them, as in an exam. They are carefully monitored to avoid communication or observation by other participants.

[14]A tag is clipped to the top of each questionnaire indicating the student's registration number. Students are instructed to present their tag in class for payment. Each cash payment is inserted in an envelope with the registration number written on the front; the envelope is given to the student when he or she presents the corresponding tag. Payment after the experiment (one week later in our case) is standard in classroom experiments and is common to all of our treatments. We note that pre-payment vs. post-payment may affect behavior (Rosenboim & Shavit, 2012), potentially due to endowment effects that can make subjects more selfish (List, 2007).

and selfishness in the dictator games (direct and delegated). An untruthful decision is indicated by a subject's choice of the untruthful message in the direct treatment and by choice of the untruthful Agent in the delegated treatment. Likewise, a selfish decision in the dictator game is indicated by choice of the unequal split Option B in the direct treatment and by choice of the "selfish" Agent B in the delegated treatment.

Our baseline (direct treatment) results are broadly consistent with prior literature on dichotomous choice experiments. 25.3 percent of subjects are untruthful in the direct deception game and 43.2 percent of subjects are selfish in the direct dictator game.[15] As in prior work (Gneezy, 2005; Hurkens & Kartik, 2009), lie aversion leads to fewer "selfish" (unequal) splits in the direct deception game – where the unequal split is associated with a lie – than in the direct dictator game.

Under delegation, the propensity to lie is 51.1 percent, 25.8 percentage points higher than in the direct treatment. In the delegated dictator game, the proportion of subjects choosing the selfish Agent, 48.3 %, is almost identical to the proportion choosing the lying Agent in the delegated deception treatment. It thus appears that delegation eliminates lie aversion in our experiment, although we offer an alternative interpretation in a moment.

[15]While an exact parallel to our payoffs is hard to find in the literature, our percentages are roughly in line with other work dating back to the original dichotomous choice experiments of Kahneman, Knetsch & Thaler (1986). For example, in their baseline dictator games, Fershtman, Gneezy & List (2012) find that 27.5% of U.C. San Diego subjects choose a "selfish" ($11/$2) versus "fair" ($8/$8) allocation, and Dana et al. (2007) find that 26% of their University of Pittsburgh dictators choose selfish ($6/$1) versus fair ($5,$5). In both cases, a "selfish" choice involves a welfare cost and a larger loss to the Receiver than in our experiment, which may explain the higher (43%) rate for choice of the unequal-split Option in our dictator game. On the deception side, Gneezy (2005) finds that 36% of his Israeli subjects are untruthful with weaker monetary incentives to lie than in our experiment and 52% are untruthful with stronger monetary incentives to lie. Roughly 42% of subjects are untruthful in the Arizona and California control experiments of Innes & Mitra (2013), also using the Gneezy (2005) design. In our case, the statements defining a lie are somewhat stronger than in the Gneezy (2005) game and the baseline (generous) option is an equal split, versus an option that benefits the Receiver relatively more. These differences may explain why a lower proportion of our subjects – roughly 25% – lie in the direct deception treatment.

Consistent with prior work on delegation in dictator games, a larger share of delegating dictators choose the selfish Agent (48.3%) than choose to be selfish in the direct dictator game (43.2%), although the difference is not statistically significant. Our design is likely to produce smaller effects of delegation on dictator decisions than found in prior work. [16] The reason is that, unlike prior delegation experiments, our Senders are confronted with precise information about what their chosen Agent does on their behalf. Hamman et al. (2010, p. 1844) write that in order to "reintroduce the social pressure or obligation to behave altruistically," principals could be informed of their agents' decisions and required to certify or override them, exactly as they do in our experiment. Our dictator results may be a symptom of this prescription.

As indicated in Table 1, average behavior in the different treatments masks heterogeneity in responses. For example, in the delegated deception game, the overall proportion of liars is 51 percent. However, when decomposed between men and women, rates of untruthfulness are significantly different; 61.4 percent of men choose the lying agent, whereas 42.3 percent of women do so; the difference is statistically significant (z=1.89). Similarly, in the delegated dictator game, 58.7% of men are selfish, compared with 36.6 percent of women; again the difference is significant (z=2.11). We thus find that gender differences in preferences persist under a delegated choice environment, in both deception and dictator games. [17]

[16] Due to different experimental designs, we are reluctant to make direct comparisons between our dictator results and those from prior literature. This said, Bartling & Fischbacher (2012) find that the total rate of unfairness is a statistically insignificant 7.6 percentage points higher under delegation (59.1 percent) than it is with no delegation (51.5 percent) in their one-shot games with more extreme payoff choices than ours (z=0.90).

[17] For a study on gender effects in direct dictator games, see Eckel & Grossman (1998); on gender effects in direct deception games, see Dreber & Johannesson (2008) and Friesen & Gangadharan (2012). These studies indicate that women tend to be more generous and less deceitful than men, as in our experiment. The behavioral differences across genders indicate both that active decisions are made in our games and that, to the extent there is any random behavior, it does not bias in favor of our inferences. For example, for the men (the majority of our black lie sample), any randomization would bias against our findings by favoring a higher (closer to 50%) proportion of untruthful subjects in the direct treatment and a lower (closer to 50%) proportion of untruthful subjects in the delegated treatment.

Across the genders, the effects of the delegation (vs. direct) treatment are similar. For men, the propensity to lie is 30 percent higher under delegation; for women, delegation increases the propensity to lie by 21.8 percent. Corresponding changes in the dictator game are almost identical between the genders, with delegation raising rates of selfishness by 5.8 percent for men and 6.9 percent for women. We therefore turn to our main hypotheses using our entire sample of both genders.

Following our hypotheses in Section 2, we gauge the statistical significance of the delegation effect on lie aversion in several ways. First is the raw difference between untruthfulness in the delegated and direct deception games (Hypothesis 1). The test statistic for the null hypothesis of zero difference is 3.71 (p-value < 0.0001). Second is a difference-in-difference test (delegated minus direct, deception minus dictator) that controls for potential effects of delegation on preferences over the payoff allocations by netting out corresponding changes in dictator choices. Consistent with Hypothesis 2, the difference-in-difference is 20.7 percentage points; delegation raises the propensity to lie by over 20 percentage points more than it raises selfishness. The test statistic for the null hypothesis of zero difference-in-difference is 2.02, significant at the (two-sided) 5 percent level.

Third is a Hurkens-Kartik (2009) statistic for the hypothesis that the proportion of the "selfish" who are untruthful is higher under delegation than under direct choice (Hypothesis 3). For a between-subjects design like Experiment 1, the corresponding null hypothesis is that the ratio of proportions - the fraction of subjects who are untruthful in the deception game divided by the fraction of subjects who are "selfish" in the dictator game – is the same in the delegation and direct treatments. In the Appendix, we derive a test statistic that is approximately distributed as a standard normal under this null.[18] For our samples, this

[18]The statistic is slightly different than constructed by Hurkens & Kartik (2009), who assume that the true fraction of subjects preferring the selfish allocation is given by the corresponding sample proportion.

statistic is 2.242 (p-value 0.028), consistent with Hypothesis 3. [19] We conclude that, for our payoffs, the delegation treatment leads to a significantly higher propensity for lies among the selfish.

Fourth is a test for whether delegation reduces truthfulness *proportionately* more than it reduces generosity. Table 1C reports that the ratio of truthfulness under delegation to that under the direct deception treatment is 0.655. The corresponding ratio in the dictator game is 0.910. Delegation proportionately reduces the rate of truthfulness by an estimated 25.5% more than it proportionately reduces the rate of generosity (z =1.78).[20]

Panel D of Table 1 confirms our conclusions with regressions that control for course effects and gender. The regressions include data from all four treatments. The endogenous variable is the zero-one indicator for an untruthful or selfish choice in, respectively, a deception and dictator game. The key regressors are the dummies for deception, delegation, and their interaction. The deception dummy indicates whether the game is deception (=1) or dictator (=0). Similarly, the delegation dummy indicates whether the treatment is delegation (=1) or direct (=0). The interaction is our main interest, capturing the difference-in-difference effect of delegation on untruthfulness vs. selfishness. Gender-treatment interactions are added in Model 2. Models 1 and 2 are linear probability, and Model 3 is a Probit estimation for our preferred specification.

In all models, the difference-in-difference interaction is statistically significant (at five percent). The coefficients are remarkably stable across models. The estimated delegation

The statistic derived in the Appendix is based on random sampling from all four distributions, deception and dictator, delegated and direct.

[19] In the delegated samples, the ratio of untruthfulness to selfishness, (0.5106/0.4828), is greater than one, although the difference is not statistically significant. For this case, a potential alternative Hurkens-Kartik test would be for the null that the corresponding ratio for the direct samples equals one. The test statistic for this alternative null is 2.54 (p=0.013), giving rise to the same conclusion. To determine p values, we use a conservative degrees of freedom equal to the minimum treatment sample size minus one (df=86).

[20] We thank an anonymous referee for suggesting this test. The test statistic is constructed using the same approach as for the Hurkens-Kartik statistic (see Appendix).

effect on untruthfulness, over and beyond its effect on the propensity to be selfish, is an increase of between 20.8 percentage points (Model 2) and 24.2 percentage points (Model 3). Male gender has a significantly positive effect on the propensity to lie and be selfish, but there is no evidence of significant differences in the gender effect across games and treatments.

The propensity to lie is significantly lower than the propensity to be selfish, as indicated by coefficients on the Deception dummy. These coefficients confirm prior work documenting lie aversion in direct games. However, we cannot reject the hypothesis of a zero deception effect in the delegation treatments. Test statistics for the null of a zero sum of coefficients on the Deception dummy and the Deception-Delegation interaction have p-values ranging from 0.38 (Model 2) to 0.62 (Model 1).

Overall, Experiment 1 indicates that *the delegation treatment reduces lie aversion.* Under delegation, moral preferences persist in that roughly 50 percent of subjects prefer the generous (equal split) Agent in the delegated dictator game and the truthful Agent in the delegated deception game. However, unlike in the direct treatments, we find no *additional* aversion to lying under delegation, over and beyond the preference for generosity. One possible explanation for this outcome is a crowding out of moral preferences (Gneezy & Rustichini, 2000). Under delegation, subjects might evaluate one or the other moral component but not both. Experiment 2 assesses our hypotheses under a different (white lie) payoff structure without any potential for crowding out.

4 The White Lie Experiment 2

Experiment 2 mimics the design of Experiment 1 with one major difference: If the Receiver reports the incorrect color of the dot, both Sender and Receiver earn more money. Erat & Gneezy (2012) refer to this structure as a Pareto white lie situation (see also Gneezy, 2005).

Here, the specific payments are as follows: [21]

> If the Receiver reports *the true color* (say blue), then payments are $5 each to Sender and Receiver.
>
> If the Receiver reports *the incorrect color* (say green), then payments are $6 each.

In Experiment 2, there is no reason for a subject to prefer the "truthful allocation" to the "untruthful allocation" because the latter is better for both players. However, lie aversion may nonetheless prevent subjects from sending an untruthful message that solemnly swears the dot is the incorrect color. The experiment identifies the extent to which delegation reduces lie aversion, promoting the delivery of more untruthful messages.

Logistics. Using the same broad protocols as in Experiment 1, we conduct Experiment 2 in three undergraduate economics classes at U.C. Merced, with 78 Sender-Receiver pairs, 39 Senders in the direct deception game and 39 in the delegated deception game. Direct and delegated treatments are equally represented in each classroom, with equally mixed questionnaires. The proportion of males in the direct and delegated samples is 46.1% and 43.6%, respectively. In the experiment, 86 percent of queried Senders predict that their Receiver will follow their message, supporting the design intent to avoid strategic considerations.

Results from Experiment 2. Figure 4 and Table 2 present results of the experiment. One-third of subjects are untruthful in the direct treatment, compared with 56.4 percent of subjects in the delegated treatment.[22] The 23 percentage point difference is large, sta-

[21] The payoff gains to dishonesty in our experiment are in line with prior literature. With an untruthful (vs. truthful) message, our subjects earn $1 more (times the probability that the Receiver follows), compared with (for example) an average of 50 U.S. cents in Erat & Gnnezy (2012), one euro in Lopez-Perez & Spiegelman (2013), and 0.3 to 1.2 Swiss Francs (roughly 25 U.S. cents to $1.05) in Gibson et al (2013).

[22] Perhaps surprising is the similarity between rates of deception in the "black lie" Experiment 1 and the "white lie" Experiment 2. There are offsetting effects: Experiment 2 (vs. 1) gives subjects a smaller own-benefit from lying, but gives Receivers a benefit (vs. cost) from the lie. Results from Erat & Gneezy (2012) suggest that the former (own payoff) effect should favor fewer lies in Experiment 2, whereas the latter (Receiver payoff) effect should favor more lies in Experiment 2. The small difference in lie rates between the experiments likely reflects this tradeoff.

tistically significant, and strikingly similar to delegation effects in Experiment 1. Panel B of Table 2 presents regressions that control for course and gender effects, confirming the evidence from the coarse statistics.

Two conclusions emerge from Experiment 2. In the white lie context, *the delegation treatment again reduces lie aversion for a significant fraction of subjects.* In addition, Experiment 2 provides evidence that *delegation does not eliminate lie aversion.* Even with delegation, over 43 percent of subjects are truthful to their detriment, a percentage that is significantly different from zero (z-statistic 5.49, p-value < 0.001).

In Experiment 2, we do not control for preferences over allocations because of the natural presumption that all will prefer the (\$6,\$6) payoff to the truthful (\$5,\$5) counterpart. However, perhaps delegation affects the *strength* of relative preference for the former over the latter. Bartling, Fischbacher & Herz (2014) document a prevalent preference for direct (vs. delegated) decision making. Gawn & Innes (2018) find that subjects also have a relatively greater preference for direct decisions when the potentially delegated decision is "good news" vs. "bad news." For our white lie experiment – where the lie allocation is good news for the Receiver – both sets of results suggest that most subjects will prefer to directly choose the superior allocation, rather than have it chosen indirectly via an agent. If anything, our experimental results will therefore tend to *understate* the extent to which delegation reduces lie aversion; some subjects may opt to tell the truth under delegation when they would not do so directly because they gain relatively less from the superior (untruthful) allocation.

5 Robustness Experiment 3

We run an additional experiment to examine robustness of our results to a number of design features: [23]

[23] We thank anonymous referees for raising these issues.

1) Probability of Decision Implementation. In the above experiments, Message/Agent choices are implemented with two-thirds probability. Do our results persist if decisions are implemented with 100 percent probability?

2) Agent Payoff Effects. As in Hamman et al. (2010), our Agents are paid based on the number of students/principals who choose them as their Agent. Do our results persist if there are no payoff effects on Agents?

3) Strategic Effects. Sutter (2009) observes that subjects may tell the truth anticipating that their Receiver will not follow the message (sophisticated liars). Do our results persist if there are no such strategic considerations?

4) Information on Heterogeneous Decisions. In our delegation treatments, subjects are aware that Agents made both choices. Do our results persist if participants in our direct treatments also know that both decisions have been made by others?

Our follow-up experiment addresses these questions with a black lie design that stipulates the same payoffs and basic structure as described in Section 3 but with the following differences:

1) Message decisions are implemented with 100 percent probability. In order to mimic Receiver accept/reject behavior (following Gneezy, 2005), each dictator's allocation choice is implemented with 80 percent probability (and the other Option with 20% probability).

2) In the delegation treatments, students implement the decision of a prior participant who chose between (a) the two alternate Messages, GREEN or BLUE (in the delegated deception treatment) or (b) the two alternate allocations, equal or unequal (in the delegated dictator treatment). In each case, the subject chooses between two students, one of whom made one decision (e.g., Student 1 who chose a BLUE message) and a second student who

made the other decision (e.g., Student 2 who chose a GREEN message). The decision of the selected student is implemented. To ensure an absence of Agent payoff effects, subjects are informed that "your decision does not affect Students 1 and 2 in any way and they do not know about your decision."

3) Strategic effects are avoided in the deception treatments by revising the Receiver options as described in Gawn and Innes (2017). Rather than choosing whether to Report a BLUE or GREEN dot after seeing a Message, Receivers choose between two options:

> Option A: REPORT the color indicated in the Message. OR
> Option B: Make a RANDOM REPORT, meaning that we randomly select a REPORT of GREEN or a REPORT of BLUE, each with 50 percent probability.

With these choices, a decision by the Receiver *not* to follow the Message produces payments that are completely invariant to whether the Message is BLUE or GREEN. As a result, a subject chooses an untruthful or truthful Message only if he or she prefers that choice when the chosen Message is followed, regardless of the believed likelihood that the Receiver will actually follow. That is, there are no strategic effects as described in Sutter (2009). [24]

4) In the direct treatments, subjects are informed that "all possible decisions were made by at least one person" in a prior session of the experiment (Gawn & Innes, 2017).

Logistics. The experiment is run in two lower division and four upper division economics classes at U.C. Merced, enrolling a total of 323 subjects in the sender role. There are minor differences in the gender composition across the treatments; the proportion of male subjects

[24] As before, we inform subjects that "in similar experiments, most Receivers Report according to the Message." This statement is based on a prior (direct) version of this experiment in which 73.1 percent of Receivers followed the message (Gawn and Innes, 2017). The information is provided in order to convey that the Message decision is consequential, even though beliefs about Receiver behavior do not create strategic incentives. In the present experiment, 71.2 percent of "delegated" Receivers follow the Message and 74.1 percent of "direct" Receivers follow the Message.

varies from a low of 53.3 percent in the direct deception game to a high of 60.5 percent in the delegated deception game. Dot colors and option labels are randomly varied across subjects.

Results. Figure 5 and Table 3 present results of the follow-on Experiment 3. Outcomes are generally similar to those from Experiment 1. Consistent with Hypothesis 1, truthfulness declines by 32.5 percentage points when Messages are chosen via delegates than when they are chosen directly (z=4.76). Delegated decision making also promotes less generous allocations in the dictator game, by a statistically insignificant 7.8 percentage points. The difference-in-difference – the additional effect of delegation on truthfulness vs. generosity – is 24.7 percent (z=2.31), consistent with Hypothesis 2. The Hurkens-Kartik statistic indicates that the propensity for untruthfulness by the selfish rises by almost 50 percentage points in the delegated (vs. direct) treatment (z=-3.01). Similar effects are estimated in regressions that control for course effects and gender (Table 3B). We conclude that the depressing effect of delegation on lie aversion is robust to the design variations embedded in Experiment 3.

6 Discussion and Conclusion

The starting point for our study is Erat's (2013) signature result: A significant proportion of his subjects opt to delegate their decision on whether or not to lie, rather than make the decision themselves. Our experiments support Erat's (2013) primary explanation for his findings, namely, that delegation attenuates lie aversion. However, there are three crucial differences between Erat's (2013) experiment and ours: (1) By choosing to delegate in Erat (2013), a Sender cedes the decision – and therefore control – to the Agent. (2) With no parallel dictator game, effects of the delegation option on preferences over payoffs (as in Hamman et al., 2010), vs. honesty, cannot be disentangled. (3) Because there is no parallel treatment group making a deception decision without a delegation option, it is not possible to determine whether would-be liars are more likely to delegate than would-be truth tellers,

that is, whether the delegation option reduces or promotes deception overall.

Our experiments examine how exogenous delegation affects behavior, controlling for correlated effects on (i) control over outcomes, (ii) preferences over payoffs, and (iii) any desire to avoid cognitive and emotional costs of taking a decision. By thereby pinpointing the effect of delegation on lie aversion, we can address the question posed in the paper's title: does delegation promote deception? We find evidence that it does. In our experiments, subjects are more willing to lie when the lie is made with an Agent's message than when the deceptive message is sent directly by the subjects themselves.[25]

Understanding how agency relationships affect lie aversion can potentially shed light on how delegation affects moral preferences more broadly. For example, Cullen (2008) describes the case of the Canadian engineering group Acres International accused of bribing the chief executive of a Lesotho water project. Charged with violating Lesotho's anti-bribery statutes, "Acres' defence was that they were not responsible for the (bribe) payments... as these were made via an intermediary through a 'representation agreement'." Cullen (2008) writes that "there is an unwritten expectation that the 'facilitator' will do whatever is required, short of overtly criminal acts, to secure the contract or project. For a company like Acres, it is a way of distancing their bid or proposal from the unsavoury aspects of doing business in corrupt countries." Our experiments relate precisely to a situation like this in which the principal can choose an agent knowing the choices that will be made on his/her behalf.

A prominent narrative for corrupt transactions such as in the Acres' case pits reluctant

[25]The results add to the growing literature on what determines individuals' propensities for truthfulness (Rosenbaum et al., 2014). For example (and we apologize for omissions of many key papers), scholars have shown that the willingness to lie is affected by the consequences for both sides of the interaction (Gneezy, 2005), gender (Dreber & Johannesson, 2008; Friesen & Gangadharan, 2012; Houser et al., 2016), guilt aversion (Battigalli, Charness & Dufwenberg, 2013), punishment (Sanchez-Pages & Vorsatz, 2007), strategic considerations (Sutter, 2009), social cues on how often others lie (Innes & Mitra, 2013), team incentives (Conrads et al., 2013), cooperation in prior play (Ellingsen et al., 2009), and protected values (Gibson et al., 2013).

transnational corporations (TNCs) against corrupt local bureaucrats, with each company facing a prisoner's dilemma with competing TNCs in which the dominant strategy is to acquiesce in a corrupt exchange (see Collier, 2007, p.137; Dufwenberg and Spagnola, 2015). However, if delegation to an intermediary erodes the moral preferences of a principal, as we find, then causation also runs in the opposite direction: the process of intermediation promotes dishonesty and related corruption by the TNC, via an agent who is chosen precisely to be deceitful.

The attenuating effect of delegation on lie aversion has implications for the use of agents in markets, including vertical separation with outsourced suppliers and subcontractors. If delegation suppresses costly honestly, as indicated in our experiments, then firms have an added incentive to delegate, even beyond potential benefits of avoiding blame for a decision that harms another party (Bartling & Fischbacher, 2012), suppressing impulses for generosity that harm the principal's bottom line (Hamman et al., 2010), and benefits due to the information, cost, strategic and commitment considerations identified in prior literature (see, for example, Bolton & Dewatripont, 2005; Fershtman et al., 1991; Fershtman & Kalai, 1997). From an empirical point of view, these benefits are likely to be particularly relevant in economic environments where deception is normal and advantageous, such as cultures with weak moral institutions. In these corrupt environments, both private economic benefits of dishonesty to contracting firms and opportunities for contracting with dishonest agents are likely to be greater. From a normative point of view, economic effects of contractual relationships are also likely to be more pernicious than they would be absent their impacts on lie aversion.

On the flip side, the results suggest that first party interactions are likely to be more truthful than are second party interactions. For example, private party sellers of used cars

may be more honest than sales representatives in used car dealerships. This may help to explain Gneezy's (2005) survey results indicating that students overwhelmingly have this belief.

In practice, the impact of agency on dishonesty is likely to be greater than obtains here. Our Sender-Receiver experiments involve anonymous interactions with no opportunities for subsequent punishment. Audience effects (Andreoni & Bernheim, 2009) or opportunities for punishment (as in Bartling & Fischbacher, 2012; Coffman, 2011; and Oexl & Grossman, 2013) might be expected to promote even greater effects of delegation on honesty by "shifting the blame." In addition, Lazear, Malmendier & Weber (2012) show that the opportunity for exit from a dictator game leads to selection in favor of more selfish dictators. A plausible conjecture is that self-selection might also favor more dishonest agents (see Della Vigna et al., 2017; Fehrler et al., 2016). Our results suggest an additional mechanism for the promotion of dishonesty, namely, the selection of agents by principals.

Overall, the results suggest that one's sense of responsibility / ownership of a decision is important to social behavior, and to truthful communication in particular. Bartling & Fischbacher (2012) show that a measure of a subject's responsibility for a decision better explains punishment behavior of others at the receiving end than do measures of outcome-based or intention-based social preferences; how others respond to our actions depends upon how responsible they think we are for those actions (see also Engl, 2018). Our results suggest that if a subject perceives a reduced responsibility for a lie, their own moral compass is altered – even absent any objective loss of true responsibility/control and no scope for response or full understanding by those at the receiving end of the lie.

7 References

Abeler, J., A. Becker, & A. Falk. 2014. "Representative Evidence on Lying Costs." *J. Public Economics*, 113, 96-104.

Andreoni, J. & B. D. Bernheim. 2009. "Social Image and the 50-50 Norm." *Econometrica*, 77(5), 1607-36.

Bartling, B. & U. Fischbacher. 2012. "Shifting the Blame: on Delegation and Responsibility." *Review of Economic Studies*, 79 (1), 67-87.

Bartling, B., U. Fischbacher, & H. Herz. 2014. "The Intrinsic Value of Decision Rights." *Econometrica*, 82, 2005-39.

Battigalli, P., G. Charness, & M. Dufwenberg. 2013. "Deception: The Role of Guilt." *J. Economic Behavior & Organization*, 93, 227-32.

Blount, S. 1995. "When Social Outcomes Aren't Fair: The Effect of Causal Attribution on Preferences." *Org. Behavior & Human Decision Processes*, 63, 131-44.

Bolton, P., & M. Dewatripont. 2005. *Contract Theory.* MIT Press.

Charness, G. 2000. "Responsibility and Effort in an Experimental Labor Market," *J. Economic Behavior & Organization*, 42, 375-84.

Coffman, L. 2011. "Intermediation Reduces Punishment (And Reward)." *American Economic Journal: Microeconomics*, 3(4), 77-106.

Collier, P. 2007. *The Bottom Billion: Why the Poorest Countries are Failing and What Can Be Done About It.* New York. Oxford University Press.

Conrads, J., B. Irlenbusch, R.M. Rilke & G. Walkowitz. 2013 . "Lying and Team Incentives." *J. Economic Psychology*, 34, 1-7.

Cullen, R. 2008. *The Poverty of Corrupt Nations.* Toronto. Blue Butterfly Publishing.

Dana, J., D. Cain, & R. Dawes. 2006. "What You Don't Know Won't Hurt Me: Costly (But Quiet) Exit in Dictator Games." *Org. Behavior & Human Decision Processes*, 100(2), 193-201.

Dana, J., R. Weber, & J. Kuang. 2007. "Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness." *Economic Theory*, 33(1), 67-80.

Della Vigna, S., J. List, U. Malmendier, & G. Rao. 2017. "Voting to Tell Others." *Review of Economic Studies*, 84 (1), 143-81.

Dreber, A. & M. Johannesson. 2008. "Gender Differences in Deception." *Economics Letters*, 99, 197-99.

Drugov, M., J. Hamman, & D. Serra. 2014. "Intermediaries in Corruption: An Experiment." *Experimental Economics*, 17, 78-99.

Dufwenberg, M. & G. Spagnola. 2015. "Legalizing Bribe Giving." *Econ. Inquiry*, 53, 836-53.

Eckel, C. & P. Grossman. 1998. "Are Women Less Selfish Than Men? Evidence from Dictator Experiments." *Econ. Journal*, 108, 726-35.

Eisenkopf, G., & U. Fischbacher. 2015. "Naive Responses to Kind Delegation." *Managerial & Decision Economics*, 36, 487-98.

Ellingsen, T., M. Johannesson, J. Lilja, & H. Zetterqvist. 2009. "Trust and Truth." *Economic Journal*, 119, 252-76.

Engl, F. 2018. "A Theory of Causal Responsibility Attribution." Working Paper, University of Cologne.

Erat, S. 2013. "Avoiding Lying: The Case of Delegated Deception." *J. Economic Behavior & Organization*, 93, 273-78.

Erat, S. & U. Gneezy. 2012. "White Lies." *Management Science*, 58(4), 723-33.

Falk, A., & M. Kosfeld. 2006. "The Hidden Costs of Control." *Amer. Econ. Rev.*, 96(5), 1611-30.

Falk, A., & N. Szech. 2013. "Morals and Markets." *Science*, 340, 707-11.

Fehrler, S., U. Fischbacher & M. Schneider. 2016. "Who Runs? Honesty and Self-Selection into Politics." Working Paper, IZA.

Fershtman, C., & U. Gneezy. 2001. "Strategic Delegation: An Experiment." *RAND J. Economics*, 32(2), 352-68.

Fershtman, C., U. Gneezy, & J. List. 2012. "Equity Aversion: Social Norms and the Desire to be Ahead." *American Economic Journal: Microeconomics*, 4(4), 131-44.

Fershtman, C., K. L. Judd, & E. Kalai. 1991. "Observable Contracts: Strategic Delegation and Cooperation." *Intern. Econ. Rev.*, 32, 551-59.

Fershtman, C. & E. Kalai. 1997. "Unobserved Delegation." *Intern. Econ. Rev.*, 38, 763-74.

Fischbacher, U. & F. Föllmi-Heusi. 2013. "Lies in Disguise - An Experimental Study on Cheating." *J. European Economic Association*, 11(3), 525-47.

Friedman, E. 2011. "Romney Explains Record of Hiring Illegal Immigrants as Lawn Keepers." *ABC News*, October 18.

Friesen, L. & L. Gangadharan. 2012. "Individual Level Evidence of Dishonesty and the Gender Effect." *Economics Letters*, 117, 624-626.

Gawn, G. & R. Innes. 2017. "Pure Lie Aversion in Two-Sided Games: Another Look." Working Paper, U.C. Merced.

Gawn, G. & R. Innes. 2018. "Delegating Selfish vs. Generous Decisions: An Experimental Test of the Machiavelli Conjecture." Working Paper, U.C. Merced.

Gibson, R., C. Tanner & A.F. Wagner. 2013. "Preferences for Truthfulness: Heterogeneity among and within Individuals." *Amer. Econ. Rev.*, 103(1), 532-48.

Gneezy, U. 2005. "Deception: The Role of Consequences." *Amer. Econ. Rev.*, 95, 384-94.

Gneezy, U. & A. Rustichini. 2000. "A Fine is a Price." *J. Legal Studies*, 29, 1-17.

Gurdal, M., J. Miller, & A. Rustichini. 2013. "Why Blame?" *J. Political Econ.*, 121(6), 1205-35.

Haisley, E. & R. Weber. 2010. "Self-Serving Interpretations of Ambiguity in Other-Regarding Behavior." *Games & Economic Behavior*, 68(2), 614-25.

Hamman, J., G. Loewenstein, & R. Weber. 2010. "Self-Interest Through Delegation: An Additional Rationale for the Principal-Agent Relationship." *Amer. Econ. Rev.*, 100(4), 1826-46.

Houser, D., J. List, M. Piovesan, A. Samek, & R. Weber. 2016. "Dishonesty: From Parents to Children." *European Econ. Rev.*, 82, 242-54.

Hurkens, S. & N. Kartik. 2009. "Would I Lie to You? On Social Preferences and Lying Aversion." *Experimental Economics*, 12(2), 180-92.

Innes, R. & A. Mitra. 2013. "Is Dishonesty Contagious?" *Econ. Inquiry*, 51(1), 722-34.

Kahneman, D., J. Knetsch, & R. Thaler. 1986. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *Amer. Econ. Rev.*, 76(4), 728-41.

Lazear, E., U. Malmendier, & R. Weber. 2012. "Sorting in Experiments with Application to Social Preferences." *American Economic Journal: Applied Economics*, 4(1), 136-63.

Levitt, S. & J. List. 2007. "What Do Laboratory Experimets Measuring Social Preferences Reveal About the Real World?" *J. Economic Perspectives*, 21(2), 153-74.

List, J. 2007. "On the Interpretation of Giving in Dictator Games." *Journal of Political Economy*, 115 (3), 482-493.

Lopez-Perez, R. & E. Spiegelman. 2013. "Why do People Tell the Truth? Experimental Evidence for Pure Lie Aversion." *Experimental Economics*, 16(3), 233-47.

Machiavelli, N. 2003. *The Prince*. London: Penguin Classics.

McLean, B., & P. Elkind. 2003. *The Smartest Guys in the Room: The Amazing Rise and Scandalous Fall of Enron*. New York: Portfolio.

Nelkin, D. 2004. "Moral Luck." In: *The Stanford Encyclopedia of Philosophy*, E. Zalta, ed., Stanford University.

Oexl, R. & Z. Grossman. 2013. "Shifting the Blame to a Powerless Intermediary." *Experimental Economics*, 16, 306-12.

O'Rourke, D. 1997. *Smoke from a Hired Gun: A Critique of Nike's Labor and Environmental Auditing in Vietnam.* Transnational Resource & Action Center. San Francisco.

Rosenbaum, S., S. Billinger & N. Stieglitz. 2014. "Let's be Honest: A Review of Experimental Evidence of Honesty and Truth-Telling." *J. Economic Psychology*, 45, 181-96.

Rosenboim, M. & T. Shavit. 2012. "Whose Money Is It Anyway? Using Prepaid Incentives in Experimental Economics to Create a Natural Environment." *Experimental Economics*, 15, 145-57.

Sánchez-Pagés, S. & M. Vorsatz. 2007. "An Experimental Study of Truth-telling in a Sender-Receiver Game." *Games & Economic Behavior*, 61, 86-112.

Sutter, M. 2009. "Deception through Telling the Truth? Experimental Evidence from Individuals and Teams." *Economic Journal*, 119, 4760.

60 Minutes. 2015. "Lumber Liquidators Linked to Health and Safety Violations." *CBS News*, March 1.

Appendix: The Hurkens-Kartik Statistic for Experiment 1

Let p_k = mean of Bernoulli distribution k, q_k = sample mean from k, and n_k = number of observations for sample k. We have four independent random samples, $k = 1$ (delegated deception), $k = 2$ (direct deception), $k = 3$ (delegated dictator), and $k = 4$ (direct dictator). Each Bernoulli observation takes a value of one if the choice is untruthful (selfish), and zero if the choice is truthful (generous). The null hypothesis is that the ratio of untruthful to selfish for delegated decisions equals the ratio of untruthful to selfish for direct decisions:

$$H_0 : (p_1/p_3) - (p_2/p_4) = 0 \longleftrightarrow (p_1 p_4) - (p_2 p_3) = 0; \qquad H_1 : (p_1 p_4) - (p_2 p_3) \neq 0$$

A few preliminary observations aid derivation of the test statistic for this null.

(i) $E(q_k q_m) = p_k p_m$ and $E(q_k^2 q_m^2) = E(q_k^2)E(q_m^2)$ for $k \neq m$ by independence of observations;

(ii) $E(q_k^2) = E\{[\sum_i x_i^2 + \sum_i \sum_{j \neq i} x_i x_j]/n_k^2\} = (1/n_k)[p_k + (n_k - 1)p_k^2]$ where x_i and x_j are (Bernoulli) observations from sample k and the second equality follows from independence of observations ($E(x_i x_j) = p_k^2, i \neq j$) and $E(x_i^2) = p_k$; (iii) for $k \neq m$, using properties (i)-(ii),

$$V(q_k q_m) = E\{(q_k q_m - E(q_k q_m))^2\} = E\{(q_k q_m)^2\} - E\{q_k q_m\}^2 = E(q_k^2) \cdot E(q_m^2) - p_k^2 p_m^2$$
$$= p_k p_m \{1 + (n_k - 1)p_k + (n_m - 1)p_m - (n_k + n_m - 1)p_k p_m\}/(n_k n_m)$$

(iv) by independence of samples, the variance of the difference in sample mean products is: $V(q_1 q_4 - q_2 q_3) = V(q_1 q_4) + V(q_2 q_3)$.

By these properties, the Central Limit Theorem and the Law of Large Numbers, the following statistic is approximately distributed as a standard normal under the null:

$$t_0 = \{(q_1 q_4) - (q_2 q_3)\}/\{[W(q_1 q_4) + W(q_2 q_3)]^{1/2}\} \overset{a}{\sim} N(0,1) \text{ under } H_0$$

where W = estimated $V = q_k q_m \{1 + (n_k - 1)q_k + (n_m - 1)q_m - (n_k + n_m - 1)q_k q_m\}/(n_k n_m)$ for $(k = 1, m = 4)$ and $(k = 2, m = 3)$.

Figure 1. Deception Game Tree for Experiment 1 (When Blue is True)

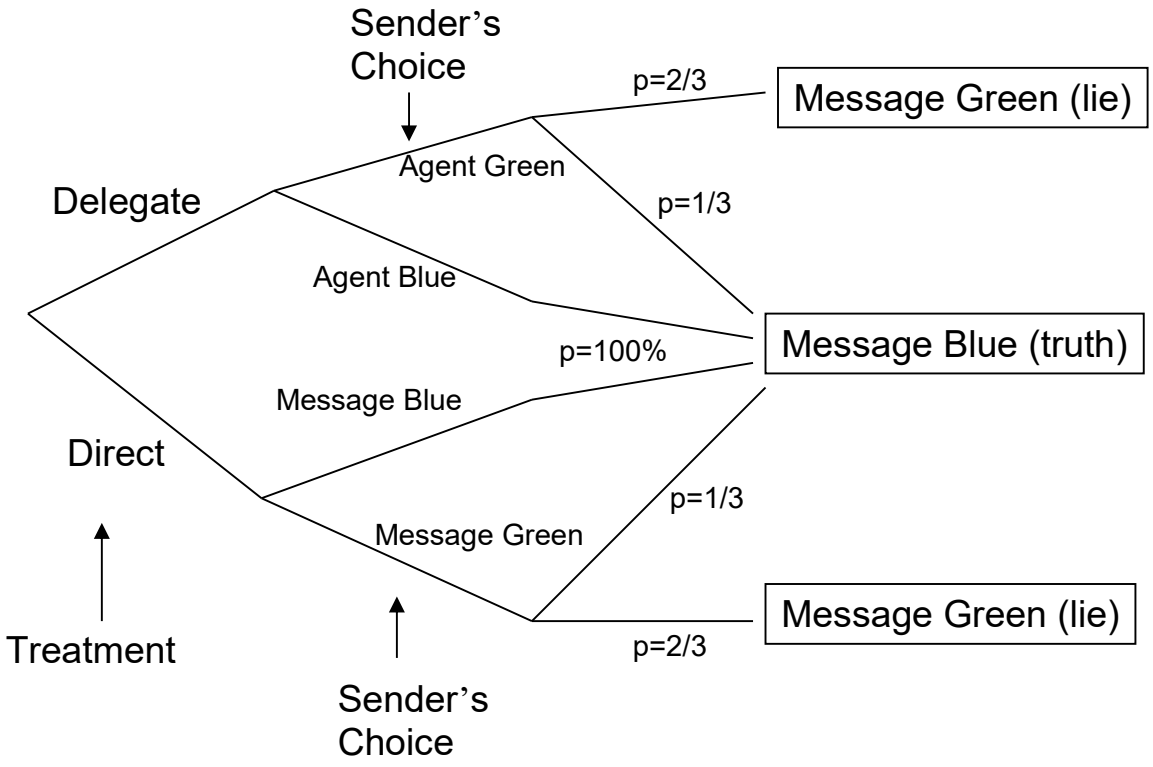

Figure 2. Dictator Game Tree for Experiment 1

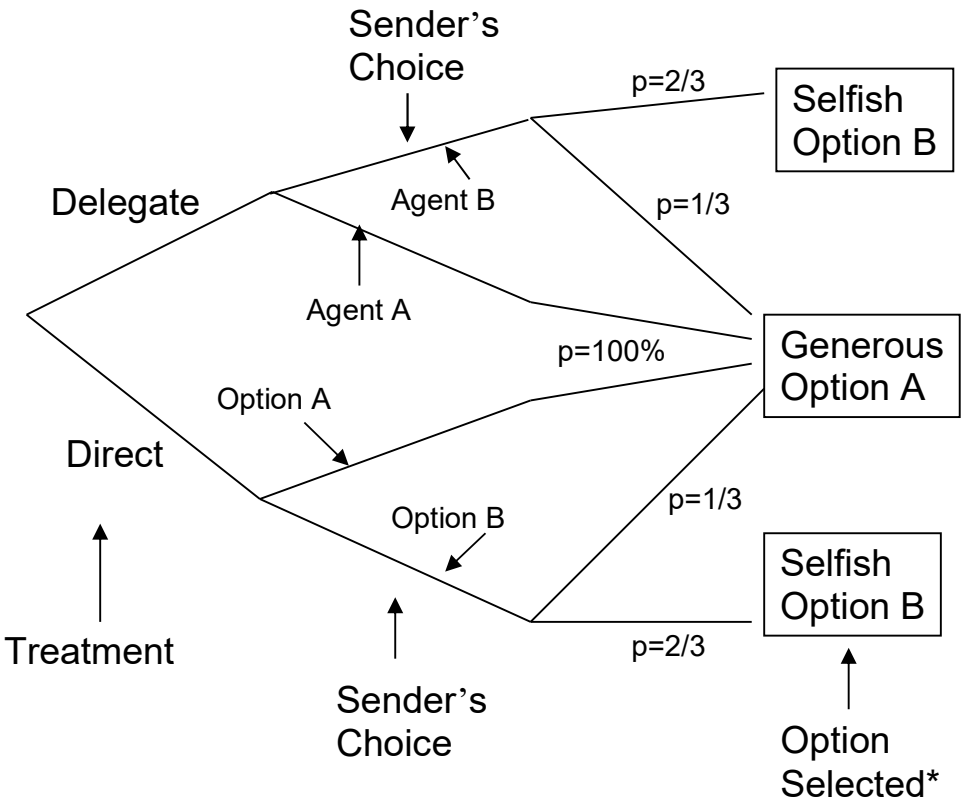

*The selected Option is implemented with 80 percent probability, and the other Option is implemented with 20 percent probability.

Figure 3. Black Lie Experiment 1 Results

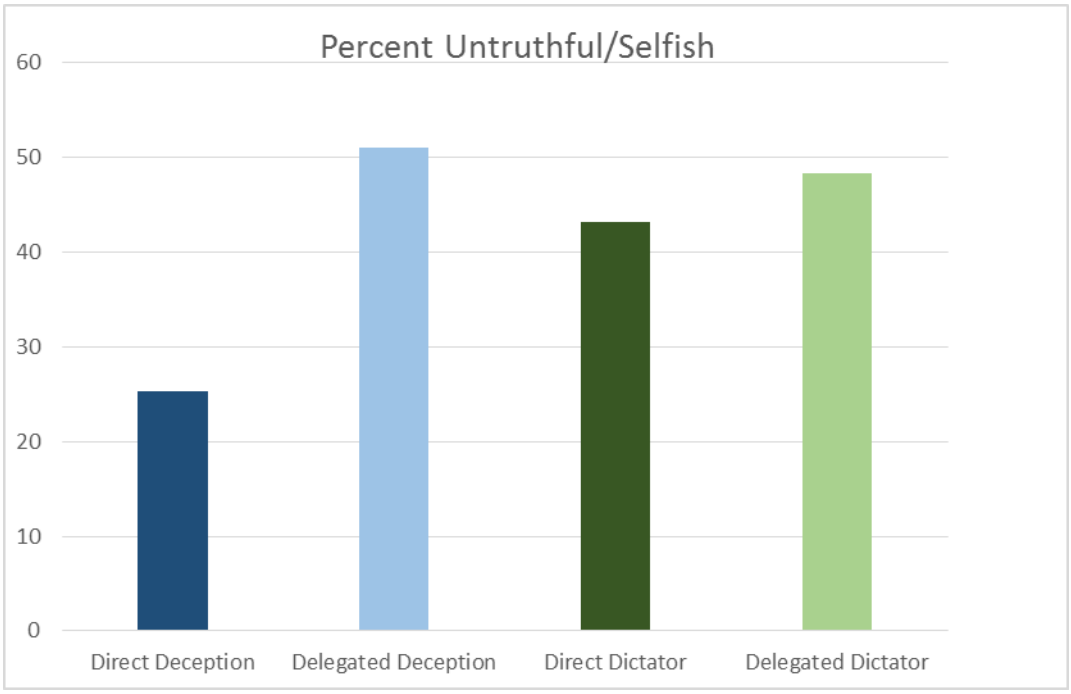

Figure 4. White Lie Experiment 2 Results

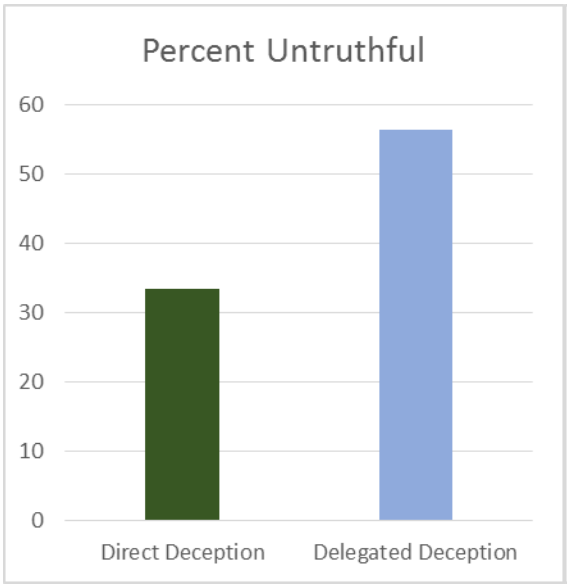

Figure 5. Black Lie Robustness Experiment 3

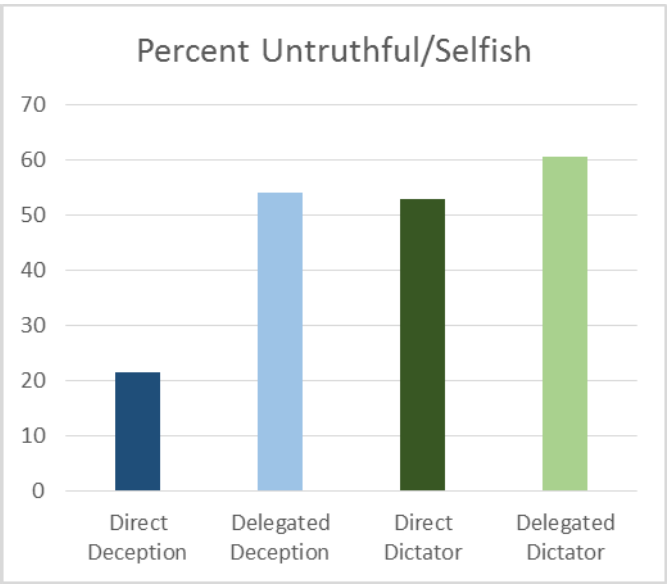

Table 1. The Black Lie Experiment 1

A. Summary and Difference Statistics

Percent Untruthful/ Selfish ↓	Sample ↓ Full	Male	Female
Delegated Deception (% Unruthful)	51.06% (N=94)	60.4% (N=48)	41.3% (N=46)
Direct Deception (% Unruthful)	25.29% (N=87)	30.4% (N=46)	19.5% (N=41)
Delegated Dictator (% Selfish)	48.28% (N=87)	58.7% (N=46)	36.6% (N=41)
Direct Dictator (% Selfish)	43.18% (N=88)	52.9% (N=51)	29.7% (N=37)
Difference (z-statistic) ↓			
Delegated Deception - Direct Deception	25.77% (3.71)***	30.0% (3.06)***	21.8% (2.28)**
Delegated Dictator - Direct Dictator	5.10% (0.68)	5.8% (0.38)	6.9% (0.70)
Difference-in-Difference	20.67% (2.02)**	24.2% (1.72)*	14.9% (1.04)

B. Hurkens-Kartik Statistics: Percent Untruthful / Percent Selfish

	Delegated	Direct	Difference (z-statistic)
Difference in Ratios: % Untruthful / % Selfish	1.0576	0.5857	0.4719 (2.24)**

C. Proportionate Effect of Delegation on Truthfulness/Generosity

	Dictator (Generosity)	Deception (Truth)	Difference (z-statistic)
Difference in Ratios: %Truth (%Generous) (Delegated) ÷ %Truth (%Generous) (Direct)	0.91030	0.6550	0.2553 (1.78)*

Notes: *,**,*** denotes significant at 10%, 5%, 1% (two tail).

Table 1. The Black Lie Experiment 1 (continued)

D. Regressions

Dependent Variable: Untruth/Selfish
N=356

	Model		
	(1) OLS[A] Marginal Effect (Standard Error)	(2) OLS[A] Marginal Effect (Standard Error)	(3) Probit[B] Marginal Effect (z statistic)
Deception Dummy	-0.1733 (0.0685)**	-0.1296 (0.0871)	-0.2000 (-2.579)**
Delegate Dummy	0.0596 (0.0726)	0.0506 (0.0918)	0.0592 (0.785)
Deception *Delegate	**0.2101 (0.0994)****	**0.2079 (0.0998)****	**0.2417 (2.224)****
Course Effects	Yes	Yes	Yes
Gender (Male)	0.1823 (0.0514)***	0.2128 (0.0911)**	0.1911 (3.441)***
Male*Deception	No	-0.0795 (0.1006)	No
Male*Delegation	No	0.0200 (0.1003)	No
R^2	0.1200	0.1217	0.0945

Notes:

[A] OLS regressions with robust standard errors. [B] Probit MEs evaluated at means, robust standard errors. **,*** denotes significant at 5%, 1% (two tail).

Table 2. The White Lie Experiment 2

A. Summary and Difference Statistics

	Percent Untruthful	Number of Observations
Delegated Sample	56.41%	39
Direct Sample	33.33%	39
Difference (z-statistic)	23.08% (2.11)**	

B. Regressions

Dependent Variable: Untruth
N=78

	(1) OLS[A] Marginal Effect (Standard Error)	(2) OLS[A] Marginal Effect (Standard Error)	(3) Probit[B] Marginal Effect (z statistic)
Delegation Dummy	**0.2308** **(0.1110)****	**0.2343** **(0.1105)****	**0.2454** **(2.160)****
Gender (Male)	No	0.1388 (0.1100)	0.1522 (1.310)
Course Effects	Yes	Yes	Yes
R^2	0.0770	0.0954	0.0730

Notes:

[A] OLS regressions with robust standard errors. [B] Probit MEs evaluated at means, robust standard errors.

,* denotes significant at 5%, 1% (two tail).

Table 3. The Black Lie Robustness Experiment 3

A. Summary and Difference Statistics

Treatment ↓ →	Delegated	Direct	Difference (z-statistic)
Deception (% Unruthful)	54.02% (N=87)	21.51% (N=93)	32.51% (4.76)***
Dictator (% Selfish)	60.56% (N=71)	52.78% (N=72)	7.78% (0.94)
Difference (z-statistic)	-6.54% (-0.83)	-31.27% (-4.30)***	24.72% (2.31)**
Hurkens-Kartik: (% Untruth / % Selfish)	0.8920	0.4075	0.4845 (3.014)***

B. Regressions

Dependent Variable: Untruth/Selfish
N=320

	Model		
	(1) OLS Marginal Effect (Standard Error)	(2) OLS Marginal Effect (Standard Error)	(3) Probit Marginal Effect (z statistic)
Deception Dummy	-0.3138 (0.0754)***	-0.3417 (0.0942)***	-0.3337 (-4.080)***
Delegate Dummy	0.0717 (0.0854)	0.0853 (0.1062)	0.0724 (0.853)
Deception *Delegate	**0.2444 (0.1093)****	**0.2414 (0.1100)****	**0.2691 (2.353)****
R^2	0.1062	0.1070	0.0813

Notes: Robust standard errors. Probit ME's evaluated at means. All models include course effects and Male gender. Model (2) also includes interactions between the Male dummy and (a) the Deception dummy and (b) the Delegate dummy, respectively. **,*** denotes significant at 5%, 1% (two tail).