Pure Lie Aversion in Two-Sided Games: Another Look

Glynis Gawn Robert Innes *

November 2017

Abstract

In a recent paper, Vanberg (2017) presents experimental evidence suggesting that the "pure lie aversion" identified in a famous paper by Erat and Gneezy (2012) can instead be explained by strategic motives: anticipating that a Receiver will not believe their communication, Senders tell the truth only because it is in their monetary self-interest to do so, not because they are averse to the alternative of lying. Vanberg (2017) presents a new experiment in which there are no strategic motives and players always lie when they benefit materially from doing so. In the present paper, we replicate Vanberg's new experiment, but add an interpretive survey question revealing that subjects perceive a new moral motive to lie that is not present in Erat and Gneezy (2012). We go on to present new Sender-Receiver experiments in which Senders have no strategic motives for their message decisions and in which the new moral motive of Vanberg (2017) is not present. We find that a large proportion of subjects tell the truth despite monetary incentives to lie, lending further support to Erat and Gneezy's (2012) original interpretation of their findings.

Keywords: Deception, Lie Aversion, Moral Preferences

JEL Classifications: D03

*Economics, School of Social Sciences, Humanities & Arts, U.C. Merced, 5200 N. Lake Rd., Merced, CA 95343. Emails: ggawn@ucmerced.edu, rinnes@ukcmerced.edu. Phone: (209) 228-4872.

1 Introduction

Recent experiments reveal that a substantial proportion of people, in a variety of circumstances, are unwilling to lie to achieve a superior monetary outcome for themselves. Variations on this conclusion stem from original work of Gneezy (2005), who finds that the preference for truthfulness is affected by the size of the gain to the Sender (transmitter of the lie) and the size of the loss to the Receiver (recipient of the lie). Many other factors have also been shown to affect the propensity for truthfulness, including (for example) the size of the lie (Lundquist et al., 2009; Fischbacher and Heusi, 2013; Gneezy et al., 2013; Gneezy et al., 2016), gender (Dreber & Johannesson, 2008; Friesen & Gangadharan, 2012), social cues (Innes & Mitra, 2013), team incentives (Conrads et al., 2013), cooperation in prior play (Ellingsen et al., 2009), punishment (Sanchez-Pages & Vorsatz, 2007), and individual value systems (Gibson et al., 2013). [1]

A central issue in much of this research is the role of other-regarding preferences. A vast literature documents that individuals care about what others receive. So when a lie benefits the teller and harms the recipient, other-regarding preferences come into play. An individual may avoid a lie because it harms the other player, not because the lie itself imparts a moral cost. Distinguishing the lie aversion component of a decision is a focus of the early deception literature studying "black lies" that can hurt the recipient if believed (Gneezy, 2005; Hurkens and Kartik, 2009).

In a key paper, Erat and Gneezy (EG, 2012) avoid this potential confound altogether by posing a "Pareto white lie" situation in which both the Sender and the Receiver benefit if a lie is told and believed. In such cases, only a pure moral cost of a lie to the Sender can motivate truthfulness. Their experiment consists of a Sender transmitting a message to a Receiver about the outcome of a die roll, after which the Receiver is asked to report the true roll, based only on the Sender's message. Both Sender and Receiver earn more money if the Receiver reports a different number than the true roll. EG find that a large proportion of Senders are truthful even in the Pareto white lie situation, powerful evidence for pure lie

[1] See Rosenbaum et al. (2014) for a recent survey. We apologize for omission of many other important studies on dishonesty and deception.

aversion.

A recent paper by Vanberg (2017) disputes EG's conclusion, arguing that truth telling can instead be explained by the strategic considerations first identified by Sutter (2009): Anticipating that a Receiver will not follow their message (and will instead report a different die roll than the Sender indicates), Senders tell the truth only in order to advance their own monetary self-interest. In an experiment replicating EG's Pareto white lie design, Vanberg (2017) finds that a large majority of truth tellers expect their Receiver not to follow their message and make their message decision in order to maximize their own payoff, rather than due to lie aversion per se.

Vanberg designs a new Pareto white lie experiment in which strategic considerations are eliminated by a variety of modifications designed to ensure that Receivers always follow their respective Sender messages. In the new experiment, not a single Sender tells the truth when it is materially disadvantageous to do so. That is, Vanberg finds no evidence of pure lie aversion.

In this paper, we present a series of experiments that (1) suggest an alternative interpretation of Vanberg's (2017) new experiment, and (2) address Vanberg's critique of the EG experiments but otherwise preserve key qualitative features of the EG design. Our first experiment replicates Vanberg's new experiment, but adds an interpretive question indicating how subjects view an untruthful message. The answers reveal that the Vanberg design introduces a new moral motive to lie: advancing the collective interests of the "Team" comprised of Sender and Receiver. Subjects generally do not view an untruthful message as a "lie," but rather as "the right thing to do" because it "helps the Team."

In our second and third experiments, we revise the Gneezy (2005)/ EG design so as to preclude any strategic incentives, but without adding the features of Vanberg's new design that promote the Team identity. In these experiments, a Sender's belief about the likelihood that the Receiver will follow his or her message is completely irrelevant to the sign of the net material benefit from a lie vs. a truth, as in a clever paper by Gneezy et al. (2013). [2]

[2]Our approach differs from Gneezy et al. (2013) primarily by its focus on situations that control for other

Hence, in a Pareto white lie situation (Experiment 2), truth telling can only be motivated by pure lie aversion. Similarly, in a black lie situation (Experiment 3), only lie aversion can explain a greater propensity for truthfulness than for generosity in a parallel dictator game. In both cases, we find evidence that supports the conclusions of EG and Gneezy (2005) on the pure aversion to lies.

We note that Vanberg (2017) also proposes an alternate meaning of pure lie aversion, appealing to Kant (1785) and the idea that some people may never lie, regardless of the consequences. In this Kantian world, lie aversion trumps any and all other moral and monetary consequences for some people. To our knowledge, Erat and Gneezy (2012) do not confirm this Kantian notion of pure lie aversion. [3] However, they reject the null hypothesis of "no pure lie aversion" (where "pure" means separate from other moral preferences). Is this rejection invalid? This is the interpretation of Vanberg's results that we address in this paper. [4]

Section 2 describes the Vanberg (2017) critique and the corresponding purpose of, and context for, our experiments in this paper. Section 3 lays out our replication of the Vanberg experiment, followed by our new experiments in Sections 4 and 5. We conclude in Section 6.

2 The Vanberg Critique

In Erat and Gneezy (EG, 2012), each Sender sends a message to a matched Receiver indicating the outcome of a 6-sided die roll. The Sender knows the roll, and the Receiver does not. Based only on the message, the Receiver reports an outcome. The Receiver can either follow the message (by choosing the number indicated by the Sender) or not (by choosing one of the other five numbers). Payoffs are determined by the Receiver's report, and are known only by the Sender. In the "Pareto white lie" situation, payoffs are \$20 to each of

regarding preferences (e.g., white lies vs. black lies). See Section 2 for discussion.

[3] Erat and Gneezy (2012, p. 8) conclude that their "results reject the unrealistic assumption that a person's propensity to lie is independent of incentives."

[4] Vanberg (2017, p. 450) offers this interpretation at the start of his paper, indicating that his experiments "test the hypothesis that a significant part of the population avoids lying, even when doing so would result in a Pareto improvement."

the pair if the true roll is chosen by the Receiver and $30 to each if the Receiver reports a number different from the actual roll. Payments are made with 1/20 probability.[5] Receivers do not know that incentives are aligned, but have only the one cue (the message) and so tend to follow the message with high probability.

In the Erat and Gneezy (2012) experiment, 35 percent of participants tell the truth. The authors interpret this outcome as evidence of pure lie aversion "independent of social preferences for outcomes" (p. 724) and "without any confounding of distributional preferences" (p. 726) because truth-telling is materially disadvantageous to both players and can therefore only be motivated by an aversion to a lie (a dishonest message about the number rolled).

Sutter (2009) observes that, in prior deception games where Receivers make a binary choice based on a Sender message (as in the landmark game of Gneezy, 2005, for example), Senders may tell the truth anticipating that the Receiver will not follow the message. In this case, truth telling need not be motivated by lie aversion, but rather by self-interest; if the Receiver does not follow the message, then the Sender is better off materially if he or she is truthful. Erat and Gneezy (2012) argue that such a strategic motive for truth-telling is unlikely to be at play in their white lie experiment because, with the expanded message space (6 possible die rolls), a Sender must believe that there is only a tiny probability of the Receiver following (less than 1/6th) in order for a truthful message to be in the Sender's material self-interest.

Vanberg's critique is that the Sutter (2009) strategic effects are nonetheless at work in Erat and Gneezy (EG, 2012). Vanberg (2017) replicates the EG Pareto white lie experiment with 106 subjects and finds that over 30 percent of subjects tell the truth, similar to the original EG outcomes. However, in a post-experiment questionnaire designed to pinpoint motivations for this behavior, Vanberg finds that 24 of 32 truth tellers expect the Receiver not to follow their message. 19 of these 24 indicate that their message choice is made to maximize their own payoff (rather than to "always tell the truth" or to make a "purely

[5]In Erat and Gneezy (2012), one of 20 Senders is matched with one Receiver and payments made for the game played by the chosen Sender and Receiver.

random" choice).[6] The remaining 5 of the 24 motivate their choice with a desire to "always tell the truth," as do 5 of the 8 truth tellers who expect the Receiver to follow their message; the other 3 of the 8 indicate that they chose randomly. This leaves only 10 of 106 subjects (those indicating their desire to always tell the truth) who apparently tell the truth contrary to self-interest. Although this proportion is statistically significant (z= 3.32), it is small (9.4%). Vanberg (2017) concludes that his truth telling subjects are primarily motivated by self-interest due to the expectation that their respective Receivers will not follow their messages.[7]

The objective in Vanberg's (2017) new experiment is to present subjects with a Pareto white lie situation, but with no strategic incentives, that is, with no reason to expect the Receiver not to follow the Sender's message. In Vanberg's design, a Roller (Sender) observes the outcome of a die roll and sends a message to a matched Chooser (Receiver) who then, based only on the message, guesses the specific roll number. There are two cases, both played by all subjects. In one, both partners (Roller and Chooser) earn 12 euro if the Chooser reports the true roll, and 6 euro otherwise. In the other, both partners earn 12 euro if the Chooser reports an incorrect roll (not the number rolled for the Roller) and 6 euro otherwise.

To remove strategic incentives, Vanberg's experiment differs from Erat and Gneezy (2012) in a number of important respects, including:

1. A two-role design in which subjects participate in both Roller (c.f., Sender) and Chooser (c.f., Receiver) roles.
2. Fixed partners for eight rounds, in four of which one player is in the role of Roller (and their partner in the role of Chooser) and another four in which the

[6]Whether the 19 subjects assign a less than 1/6th probability of the Receiver following (so that a truth is actually in their self-interest), or do not fully understand the experiment, or are averse to a declaration to "always tell the truth" (see note 7 below), is unclear.

[7]Vanberg (2017) cautions readers on this conclusion due to the non-incentivized nature of the post-experiment questionnaire. A potential additional worry is that some of the truth telling subjects who indicate their desire to "receive the maximum payoff" may have been averse to the alternative ("in my opinion, one should always tell the truth") because of its sweeping expression of adherence to truthfulness under any circumstance, as opposed to the one circumstance faced in the experiment.

player is in the role of Chooser (and their partner in the role of Roller).

3. "Team" language with the two permanently matched players called a "Team" throughout the experiment.

4. Common incentives, with both players obtaining the same payoff, as in Erat and Gneezy's (2012) Pareto white lie treatment. Unlike EG, all players know the common incentives and the payoffs (6 and 12 euros) throughout, although the Chooser does not know whether the 12 euro payoff is associated with a truth or an untruth in any given round.

5. Choosers learn the true roll and payment values after a round is completed, so they learn whether or not their respective Roller message is in their interest.

6. Very strong incentives for the Roller to message a high-payoff (12 euro) roll number, with a 6 euro margin of benefit to both players vs. a 50 cent (U.S.) average benefit in Erat and Gneezy (2012).

All of these features "are intended to ensure that senders may be quite certain that receivers will follow their messages" (Vanberg, p. 454). Vanberg succeeds completely in this objective. Choosers in fact follow Roller messages with almost 100 percent probability and Rollers overwhelmingly predict that they will do so.[8] Moreover, every Roller in every situation chooses a message that, assuming the Chooser follows, earns the pair the higher payoff, 12 euro. All Rollers choose truthfully when the correct report produces a 12 euro payoff, and more to the point, all Rollers choose untruthfully when an incorrect report produces a 12 euro payoff.

One possible interpretation of these findings is that, once the strategic motive for truth telling is removed, there is no evidence of any lie aversion whatsoever. However, there are two potential objections to this interpretation. First, even if a subject is averse to the "lie" of messaging an incorrect die roll, he or she may still lie when it is materially advantageous to do so in order to obtain the associated monetary benefit for the two players; the lie aversion may not be strong enough to outweigh the moral benefit of increasing the partner's payoff (the other-regarding preference) and the self-interest of 6 additional euro.

[8]Only one of 144 Chooser "guesses" was different from the messaged number.

Second, while eliminating strategic incentives, Vanberg's new design cultivates a Team ethic that may provide a new moral motive to lie. The six properties of Vanberg's design (described above) satisfy prominent criteria for Team identity and cohesion in the psychology literature. Among these criteria are (see, for example, Edger, 2010; Riggio, 2011; Hackman, 2004): (1) defined interdependencies among team members with clarity about the roles that each member plays; (2) clear shared goals and outcomes; (3) stable team membership; (4) team based rewards; and (5) a focus on the collective mission. In the Vanberg experiment, the Roller is responsible for the Team, favoring the collectively advantageous course. In addition, the Roller's choice of whether or not to lie in the Team's collective interest – which could be considered a virtuous act – is observed by his or her partner, reinforcing the collective interest. Indeed, a "lie" might be considered a sacrifice for the greater good of the Team.

Purpose of Our Experiments

In our first experiment, we provide evidence on the latter conjecture by adapting the Vanberg design to a classroom setting. Our replication satisfies all of the six key properties of Vanberg's design, as described above, but adds an interpretive component. We ask subjects to indicate how they view a Roller's decision to message an incorrect roll when the Team earns more money with an incorrect guess. The answers suggest that subjects predominantly view the Team interest as the greater moral imperative in the experiment.

Our second set of experiments addresses Vanberg's principal criticism of the EG conclusions without introducing a new moral incentive to lie. We construct a simple adaptation of the EG experiment in which there is no strategic incentive to tell the truth, anticipating that the Receiver will not follow. We remove strategic incentives using an approach similar to that developed by Gneezy et al. (2013). Our experiments differ from Gneezy et al. (2013) primarily because they address potential distributional confounds (other regarding preferences) by posing a white lie situation and controlling for preferences over allocations with parallel dictator games in a black lie situation.[9] Our Senders also retain a self-interest

[9]Gneezy et al. (2013) focus on black lies without parallel dictator games. Due to our different purpose, our experiments diverge from Gneezy et al. (2013) in a number of other respects, including binary choices, stronger messages, and incentives for Senders that (like Gneezy et al., 2013) void strategic motives but

in their Receiver's response, imparting a motive to deceive that is not present in the Gneezy et al. (2013) design.

Context: What is a Lie?

The received definition of a lie may help to frame our experiments. The standard definition contains four ingredients (see Mahon (2016) for an exhaustive philosophical exposition on this subject): A lie is (1) a statement made by one who (2) believes the statement to be false, to (3) another person (4) with the intention that the other person believe that statement to be true. Add to this the simplest definition of deception (Mahon, 2016): "to cause to believe what is false."

A Gneezy (2005) lie, or an EG white lie, plausibly satisfy all four criteria. A "lie" in the Vanberg (2017) experiment likely violates both the fourth and the deception criteria. A Roller need not (likely does not) intend that an untruthful message be believed by the Chooser, only that the Chooser follows the message in the Team's collective interest. Choosers know that the Roller's message is likely to be untruthful when an incorrect guess is advantageous, and are incentivized to guess the messaged number in these cases even while not believing that the message is truthful. This departure from the conventional definition of a lie may well explain why, in our replication of the Vanberg experiment (see Section 3), our subjects view an untruthful message more as a virtuous act in the interest of the Team than as a lie.

A growing literature studies deception in one-sided situations where, other than the experimenter, there is no one at the receiving end of a lie who is affected by the communication (e.g., Fischbacher and Heusi, 2013; Abeler et al, 2014; Gibson et al., 2013; Lopez-Perez and Spiegelman, 2013). Important virtues of these designs include simplicity and an absence of interpersonal considerations and interplay. The latter voids both strategic considerations (including Sutter (2009) effects) and, presumably, confounds created by other-regarding preferences. However, one-sided deception experiments do not satisfy the third (or fourth) criterion for a lie, as the communication is not made to anyone whom the Sender could want

(unlike Gneezy et al., 2013) preserve a Sender's self-interest in Receiver responses.

to deceive. [10] In addition, as noted by Gneezy et al. (2013, p. 294), "the 'victim' of the lie is not another participant, but rather the experimenter. The identity of the victim might affect decision making." [11] In sum, "each (one and two sided) method has its pros and cons" (Gneezy et al., 2013, p. 293).

Following EG and Vanberg, our focus is on two-sided situations. Our experiments adapt EG to avoid strategic considerations but preserve the defined characteristics of a lie.

3 Experiment One: The Vanberg Design

We first adapt the Vanberg experiment to a classroom setting. Subjects are randomly assigned to one of two groups, Group A or Group B. Each student is matched to another student from the other group in a Team. The Team remains fixed throughout the experiment, but matched participants are completely anonymous.

All students make decisions in three Rounds. In the first two Rounds, decisions are made in the role of Roller (Sender). In the third Round, decisions are made in the role of Chooser (Receiver). In each of the first two Rounds, the Roller is told the outcome of a die roll (a number between 1 and 6). In each case, the Roller selects a message indicating a number rolled (whether the true roll or another number, any one between 1 and 6). Based on his or

[10] Does it matter if there is no audience whom a Sender could deceive? One might assume that a false statement must be morally worse when the statement involves deception of another person than when it does not. If so, measured lie aversion in one-sided experiments could only understate the extent of aversion to *deceptive* lies. However, the assumption is not obviously valid. In two (vs. one) sided interactions, the game context can affect the perceived morality of statements and decisions, potentially *attenuating* lie averse preferences. In a famous Harvard Business Review article, Albert Carr (1968) writes: "most bluffing in business might be regarded simply as game strategy – much like bluffing in poker, which does not reflect the morality of the bluffer."

[11] Vanberg (2017) echoes this point, suggesting that the lie aversion found in the one-sided literature may be due in part to a perception that the experimenter is losing from the lie. With two players (as in EG and Vanberg), "the subject would lie to the other participant, and thus the experimenter is more likely to be ignored" (Vanberg, p. 458). This perspective implies that "one-sided" lies do have an audience (the experimenter) – satisfying the third criterion for a lie – but that the Sender does not intend to make the audience believe the message (as there is either no reason to do so, or the experimenter knows the truth, or both). The notion that the experimenter is an audience in these experiments derives some support from recent experiments indicating subjects' concern for the appearance of honesty even in one-sided settings (Gneezy et al., 2016). However, we note that authors of one-sided experiments take measures to avoid experimenter demand effects; for example, Lopez-Perez and Spiegelman (2013, p. 238) indicate that their experiment is to learn about decision-making and subjects should "choose as you prefer."

her matched Roller's message from ONE of the first two Rounds (randomly chosen), each participant takes the role of Chooser in Round 3 and chooses a number to report.

All participants are informed from the start that incentives are aligned, that is, that the two teammates will each receive the same common payoff. The payoffs EITHER reward truth OR reward untruth. [12] In the "reward truth" case, the two teammates earn V=$12 each if the Chooser reports the actual roll number and W=$6 each if the Chooser reports an incorrect roll number (so V>W). In the "reward untruth" case, the two players earn V=$6 if the number chosen by the Chooser is the actual roll number and W=$12 otherwise (so V<W).

The two scenarios (reward truth and reward untruth) are both played in the first 2 Rounds, and the Roller chooses a message in each case. The order of the Rounds is randomly varied across participants; for some, Round 1 rewards truth and Round 2 rewards untruth; for others, Round 1 rewards untruth and Round 2 rewards truth. Rollers know the payoffs associated with each round of the experiment. Choosers do not know whether the Round (randomly) selected for their decision corresponds to the "reward truth" or the "reward untruth" payoff scenario, but know the $12/$6 common payoffs associated with each. [13] Payments are determined according to decisions made in ONE of the two possible cases (randomly selected): either when the Group A player is Roller and the Group B player is Chooser for the Roller Round randomly chosen for the Chooser's decision, or vice versa. [14]

Following Vanberg, the experiment is designed to promote the belief that Choosers will follow messages in line with Team incentives. Based on Vanberg's (2017) results (note 8), we also give participants the following information in order to reinforce this belief:

[12] We use the "reward truth" and "reward untruth" language here for expositional convenience, but do not use any language indicating "truth" or "untruth" in the instructions.

[13] When subjects are paid, they are also informed of the payments, true roll, and Roller message for the Round 3 case in which they are Chooser. As in Vanberg, Rollers are informed that Choosers will have this information after the experiment is completed.

[14] Although we indicated the mechanical details for payment as optional reading in the instructions, we did not stress these details. Rather, to avoid confusion, we simply stressed that each Team would be paid for one of the scenarios (V>W or V<W) and one of the role assignments in that scenario ("you" as Roller and the "teammate" as Chooser or vice versa.)

Note: In similar experiments, CHOOSERS almost always choose the number indicated in the message.

After completing the experiment, participants indicate their gender; rate their understanding of the experimental instructions on a scale of 1 (very clear/ well understood) to 5 (very unclear/ confusing) with a 3 expressing no opinion ("in between"); predict whether their Chooser will follow their message, with payment of an additional $1 for a correct prediction; and answer several interpretive questions about the moral content of an "untruthful" message. Participants are asked to indicate the extent to which they agree or disagree (on a 5-option scale ranging from "strongly agree" to "strongly disagree") with each of the following characterizations of a Roller's decision to message an incorrect roll number in the "reward untruth" (V<W) scenario:

a) Helping the Team;
b) A lie;
c) Enabling his/her Teammate to earn more money;
d) The right thing to do.

Logistics

An equal number of Group A and Group B questionnaires is randomly distributed to student participants. Each questionnaire contains a registration number that identifies the student for payment and, at appropriate points, a different reference number that is used to define the match in the other group. For example, a Group A questionnaire might indicate the reference number A4, meaning that the student is matched with a Group B participant who has reference number B4. In this way, students are (randomly) pre-matched across groups and a common match is defined for the full experiment. Participants are first asked to read through the instructions for the experiment and then wait for the experimenter to review the experiment verbally. After the verbal review, students proceed with the questionnaire through the first two Roller Rounds. For each Round, the questionnaire includes three questions to verify understanding and a separate decision sheet that indicates the student's reference number

and on which the student makes his or her message choice. [15] Once a student completes the two Roller decision sheets, each identified by a Group and Round indicator (for example, A1 and A2 for the Group A Round 1 and Round 2 decisions, respectively), they are asked to return these sheets to corresponding labeled boxes (A1, A2, B1 and B2) at the front of the room. They are then instructed to wait for further instructions.

After all participants have made their decisions, they are instructed to proceed to a Group-specific experimenter who has the OTHER group's Roller messages from ONE of the two Roller Rounds (randomly selected from the two, for example box A1 for Group B or box B2 for Group A). The students each show the experimenter their reference number (indicated on their questionnaire) and are given their matched player's message sheet from the randomly selected Roller Round. The students then complete the Chooser Round 3 section of the questionnaire based on their Roller's message.

The experiment is implemented with 32 participants in an upper division economics class at U.C. Merced. 67.7 percent of the sample subjects are male (with one missing observation). 25 of the 32 subjects indicate that the experiment is clear / well understood or pretty clear / fairly understood; only one subject indicates that the experiment is not very clear/ not well understood, and none indicate that the experiment is confusing. In describing results, we consider below both the full sample of 32 subjects and a restricted sample containing only the 25 subjects indicating a good understanding of the experiment. Sample instructions are provided in the online Appendix.

General Procedures (in all experiments)

In (all) our experiments, subjects are completely anonymous and assured that, following professional norms in Economics, the instructions are implemented exactly as written. Anonymity is stressed both verbally and in written instructions, with subjects told: "Your decisions are completely anonymous. No one, including us, will ever know that you are the person making the decisions you choose." Participation is purely voluntary and has no bearing on course assessment, both of which are indicated to the students before the experiment

[15] Students overwhelmingly answer the "understanding" questions correctly, with no more than one incorrect answer in each case.

begins.

Subjects' decisions are private. Decisions are made on decision sheets and are not visible to the experimenter when returned to boxes at the front of the room. Participants are seated facing forward with ample space between them, as in an exam. Students are also told not to communicate with each other for the duration of the experiment, and are monitored to ensure privacy.

Students are identified for payment based on a registration number indicated on both their questionnaire and an appended tag. They are instructed to present their tag in a designated future class session. 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. Because the experimenter does not know the payment in each envelope at the time the envelopes are collected by participants, there is no way to infer what decisions any particular individual in the experiment has made, even if the experimenter were so inclined. These procedures are fully described to the student participants at the start of each experiment.

To avoid any potential subject concern for dollar costs to the experimenters or a desire to "please" (see Lopez-Perez & Spiegelman, 2013), we 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.

Experiment 1 (Vanberg) Results

The main results of our Vanberg experiment are described in Table 1 and Figure 1. Overall, the results are consistent with those of Vanberg with some variations. Of the full sample, a large majority (24 of 32, or 75 percent) indicate a truth or an untruth whenever it is monetarily advantageous to do so. However, unlike Vanberg's subjects, a small but statistically significant fraction of our participants, 5 of 32 or 15.6 percent ($z = 2.433$), are always truthful even when it is disadvantageous to be so. Moreover, a smaller (but still significant) fraction, 12.5 percent ($z = 2.138$), are truthful in both situations even though they predict that their

teammate will follow their message.[16]

To ensure that the results do not stem from any confusion on the part of participants, we also examine a subsample of 25, those students indicating that the experiment is clear/well understood or pretty clear/fairly understood. Again, the large majority (20 of 25, or 80 percent) send the monetarily advantageous message in each situation, although a small but significant proportion, 16 percent (z = 2.182), tell the truth in both situations. A similar proportion as in the full sample, 12 percent (z = 1.846), always tell the truth and predict that their teammate will follow.[17]

Strictly speaking, these statistics lead us to reject the null hypothesis of "no lie aversion" as measured in the Vanberg context; however, the percentage of subjects who exhibit "lie averse" choices is small and, therefore, broadly consistent with Vanberg's results.[18]

The potential for conflicting moral imperatives introduced by the Vanberg framing can be seen in answers to the interpretive questions. Table 2 and Figure 2 summarize the extent to which our subjects agree or disagree with four statements characterizing a Roller's decision to send an untruthful message in the "reward untruth" scenario (V<W). In Table 2, we compare the proportion of subjects who "strongly agree" or "agree" with each statement to the corresponding proportion who "strongly disagree" or "disagree." The differences in these proportions (agree vs. disagree) are large, positive and statistically significant for the three statements indicating moral virtue for the untruthful decision, namely, "helping the Team," "enabling the Teammate to earn more money," and "the right thing to do." However, although positive, the difference is quite small (less than ten percent) and statistically insignificant for the statement describing an untruthful message as "a lie."

[16]Three of our participants choose untruthful messages in both scenarios. However, among participants indicating a good understanding of the experiment, only one student is untruthful in both cases.

[17]Among subjects predicting that the Chooser will follow (24), 16.7 percent (z=2.191) are always truthful.

[18]Surprisingly, we are somewhat less successful than Vanberg in eliminating subject beliefs that a Chooser might not follow their message. Although 75 percent of subjects (24 of 32) predict that the Chooser will follow, 25 percent do not. In Vanberg, by contrast, 93 to 94 percent of subjects expect the Chooser to follow the message. Moreover, our Chooser behavior also displays more skepticism than in Vanberg. Among our subjects, 28.1 percent choose not to follow their Roller's message (24 percent in the sample restricted to the "good understanders"). However, as indicated above, restricting attention to subjects predicting that the Chooser will follow does not alter our qualitative conclusions (note 17).

The proportion expressing agreement with the "lie" characterization (38.7 percent) is also significantly smaller than the proportion expressing agreement with the "helping the Team" descriptor (z=2.065) and "enabling the Teammate to earn more money" (z=1.880), and smaller (by almost 13 percent) but not significantly so compared with the descriptor "the right thing to do" (z=0.849). Together, these results indicate that subjects view the untruthful message less as "a lie" and more as the morally virtuous choice that promotes Team interests.

Perhaps the interpretive results could reflect self-justification motives: "liars" claim virtue in their choices only to justify them (Festinger, 1957). There are two reasons to question this claim. First, if this were the case, we would expect to see different responses for truth-tellers than for liars (Mills, 1958). Although we admittedly have a small sample of truth-tellers (5 subjects), their responses are similar to those of the liars. Three of five agree that the advantageous untruth "helps the Team" and "enables the Teammate to earn more money" (with one disagreeing in each case); two of five indicate that the untruth is "the right thing to do" (and one disagrees); and only one describes the untruth as "a lie" (and one disagrees). Second, under a Vanberg (2017) null hypothesis of no lie aversion, there is no call to self-justify; an absence of moral cost negates a compunction to rationalize the untruth.

In sum, the interpretive answers suggest that Vanberg's (2017) experiment is better understood as evidence for a competing moral imperative and not as evidence against lie aversion per se.

4 Experiment Two: White Lie

We adapt the Erat & Gneezy (2012) Pareto white lie experiment, drawing on the Lopez-Perez & Spiegelman (2013) dot framework. The experiment involves Senders and Receivers. Each Sender is anonymously matched with a Receiver from a different classroom, with student participants playing only one role (Sender OR Receiver).

Each Sender observes a dot that is either blue or green. The dot color is randomly

varied across Senders.[19] The matched Receiver does not observe the dot. The Sender selects a message that is transmitted to the Receiver. The Sender chooses between two possible messages:[20]

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

We specify "strong" messages in order to add salience to the truthfulness (vs. untruthfulness) of the Sender's choice; with the strong language, untruthful choices are clearly lies.

Based only on the message he or she receives, the Receiver chooses a Report indicating the color of the dot. The Receiver chooses between two options:

Option 1: REPORT the COLOR indicated in the Sender's Message; OR
Option 2: Make a RANDOM REPORT, meaning that we randomly select a REPORT of GREEN or REPORT of BLUE, each with 50% probability.

Depending upon the Report, the following payments are made to the Sender and Receiver:

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

Both Sender and Receiver earn more with an incorrect Receiver report, consistent with the Pareto white lie framework of EG and Vanberg. Each Receiver is told that the Sender observes the payments associated with the Report, and nothing more about payments. Senders are told that the Receiver

[19] To check for color-blindness, we ask the Sender to write down the color of the dot on their questionnaire and all Senders do so correctly.

[20] We opt for the binary choice framework (similar to Gneezy, 2005), rather than the 6-sided die framework of EG, for simplicity. Because we address strategic motives directly, we do not need an expanded message space to attenuate strategic incentives as in EG.

"will only be told that their Report determines payments and that you know the payments connected to each Report."

In addition, to indicate to Senders that their messages are likely to be believed / followed – so that their message choice is more consequential – we provide the following information:

"In similar experiments, most Receivers Report according to the Message (Option 1)."

The statement is based on a prior experiment at UC Merced where Receivers face a similar decision, with identical ("solemnly swear") messages and the same two options given here (Report according to the message or make a Random Report); in this experiment, 64 percent of Receivers follow their Sender's message rather than choose a Random Report.[21]

The principal distinction between our experiment and EG's Pareto white lie experiment is that, if Receivers "do not follow" the Sender's message (and choose a Random Report), payments are completely invariant to the Sender's message decision. As a result, if there is no lie aversion, incentives strictly favor an untruthful message regardless of the Sender's belief about the likelihood that the Receiver will follow. [22] Telling the truth can, therefore, only be motivated by a pure aversion to lies; that is, Vanberg's critique of EG is void in this experiment. In most other respects – including a one-role design, the absence of any "Team" language, Receivers who are in the dark about the payoffs / incentives and who are in separate classrooms than Senders – our design is akin to the EG experiment (but differs from Vanberg's).[23]

[21] In the prior experiment, Receivers know that the message can be either truthful to the mutual advantage of the two (Sender and Receiver) or untruthful (to the advantage of the Sender and disadvantage of the Receiver).

[22] So long as the Sender assigns a non-zero probability that the Receiver follows, incentives favor the "lie." As described below, elicited Sender beliefs about the likelihood that the Receiver follows their message overwhelmingly indicate a belief that the probability is greater than 25 percent, and only one indicates a belief less than 10 percent.

[23] The two other main differences between our experiment and EG's are the binary choice and the "solemnly swear" language. The design is similar to Gneezy et al. (2013) in that it voids strategic incentives; the Gneezy et al. (2013) design differs in that their Sender payoffs do not depend at all on the Receiver's decision, although (like here) the Receiver's payoff when not following the Sender's message is invariant to the message sent.

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 EG ($10 times the 1/20th probability of payment), 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).

At the end of the experiment, we ask Senders to (1) rate the clarity/ understanding of the experiment (as in Experiment 1), (2) report their gender and (3) reveal their belief about the likelihood that the Receiver will follow their message. Beliefs are elicited in the usual incentive-compatible way by asking subjects to indicate a predicted percentage of Receivers who will follow (by circling one of 20 five percentage point bands, 0-5, 5-10, etc.) and paying them $1 if their prediction is correct (plus or minus five percent).

The Sender-side experiment is conducted in one upper division economics course (two classes) at U.C. Merced. We obtain 54 Sender participants and follow the same general procedures as described for Experiment 1. 63 percent of the subject pool is male. Sample instructions are provided in the online Appendix.[24]

Experiment 2 (White Lie) Results

Table 3 presents results from Experiment 2 (see also Figure 3). Of the 54 participants, 22 or 40.74% tell the truth despite a monetary cost of truth to both the Sender and the Receiver. The proportion of truth tellers is significantly different from zero (z = 6.093), and slightly higher than in EG, perhaps due to the stronger language defining the message. We also examine a restricted sample containing only subjects who rate the experiment as clear/well understood or pretty clear/fairly understood. In this sample of 45 participants (83.3 percent of the full sample), 20 or 44.44% tell the truth, which again is significantly different from zero (z = 6.000). Our results confirm Erat and Gneezy's (2012) finding of pure lie aversion in a Pareto white lie experiment, despite a design that lays to rest any potential worries about a Sutter (2009) effect.

[24]The Receiver side, for both the white lie (Section 3) and black lie experiments (Section 4), is conducted in four different economics classes with no enrollment overlap. 73.1 percent of our Receivers follow the Sender messages.

Although incentives to lie persist no matter a Sender's belief about the probability that the Receiver follows, we note that the proportion of our subjects predicting large "follow" probabilities is high. Over 72 percent predict a follow rate greater than 50 percent. Moreover, interestingly, these predictions are uncorrelated with Sender message decisions. Mean predictions are 60.3 percent for "liars" and 58.2 percent for truth tellers; the proportion predicting a follow rate greater than fifty percent is 75 percent for liars and 68.2 percent for truth tellers, an insignificant difference. The lack of correlation is (arguably) surprising because the material benefit of a lie is proportional to the perceived probability that the Receiver will follow. If the moral cost of a lie is borne whenever a lie is told (vs. whenever it is told *and believed*), then one might expect to see fewer lies when perceived Receiver follow rates – and corresponding benefits of a lie – are smaller. The absence of this correlation perhaps suggests that the moral cost of a lie is only borne when the lie is thought to be believed, so that both benefits and costs of a lie are proportional to Receiver follow rates.

5 Experiment Three: Black Lie

We present a third experiment for two (related) reasons: 1) as a robustness check on the results from Experiment 2, and 2) because the moral content of a lie may well be considered different when the lie is told to promote a mutual interest (the white lie case) as opposed to a selfish and conflicting interest (the black lie case). [25]

Our black lie experiment mimics the above white lie design exactly, with the exception of the payoffs. Here, the Sender benefits from a lie, but the Receiver loses: [26]

If the Receiver report indicates *the true color* (say blue), then payments are \$5

[25] A notable difference between white lie and black lie situations is that guilt aversion can potentially explain an aversion to lies / truth-telling in the black lie case (Battigalli et al., 2013) but not in the white lie case (see EG for discusion).

[26] 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 slightly 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. As it turns out, the latter effects appear to dominate as our subjects are substantially more honest than in Innes and Mitra (2013).

each to Sender and Receiver.

If the Receiver report indicates *the incorrect color* (say green), then payments are \$7 to the Sender and \$3 to the Receiver.

Because the black lie confounds preferences over allocations (other regarding preferences) with lie aversion costs, we implement a parallel dictator game with the same payoffs (\$5/\$5 vs. \$7/\$3). To replicate Receiver follow probabilities in the dictator game, we borrow Gneezy's (2005) approach and implement dictator decisions with 80 percent probability.[27]

The design is between-subject so that each participant completes only one of the two experiments, deception or dictator. Within each of four economics classrooms, deception and dictator questionnaires are equally mixed and randomly distributed to student participants. We obtain 55 participants in the deception experiment and 52 in the parallel dictator game. The proportion of male subjects in each of the two games is slightly (but not significantly) different, with 50.9 percent male in the deception experiment and 57.7 percent male in the dictator game.

Experiment 3 (Black Lie) Results

Table 4 presents the overall results from the black lie deception and dictator experiments. Figure 3 gives a visual representation of truthfulness in the two experiments (2 and 3). Eighty percent of participants in the black lie deception game tell the truth, a higher percentage than in most other studies, most likely due to the strong (solemnly swear) language in our messages. In the dictator treatment, just over 44 percent opt for the generous (equal split) allocation. The difference between the proportion of truth-tellers (in the deception game) and of generous players (in the dictator game) is 35.8 percent and is highly significant ($z = 4.089$). Restricting the sample to those who report a good understanding of the experiment produces almost identical results; in the revised sample of 84 subjects, there is a 38.8 percent difference between rates of truthfulness and generosity ($z= 3.881$). In both cases, we find

[27]The 80 percent probability exactly mimics payoff consequences of truthful/untruthful decisions in our black lie deception game when the perceived probability of Receiver follow is 60 percent. Mean Sender beliefs (about Receiver follow rates) turn out to be slightly lower than 60 percent (at 52.8 percent), while actual Receiver follow rates are higher (at 73.1 percent). Hence, payoff matching in our dictator game is close, but not perfect. We do not believe these differences have an important effect on decisions.

evidence of "pure" lie aversion over and beyond the (other regarding) preference for the generous / equal split allocation. As in Experiment 2, the results cannot be explained by Sutter's (2009) "sophisticated lying," because truth-telling is strictly disadvantageous to the Sender, regardless of beliefs about Receiver behavior. [28]

Table 5 indicates that our black lie results are robust to controls for gender and course effects. The table presents regressions of subjects' one-zero decisions to be, respectively, truthful in the deception game and generous in the dictator game on (1) a dummy for the deception game, (2) gender, (3) course fixed effects, and (4) in our most complete models, an interaction between gender and the deception dummy. The coefficient on the deception dummy indicates the estimated excess propensity to be truthful vs. generous. In our baseline models (both linear and Probit), the estimated effect is 35.6 percent and is highly significant (z=3.93 and z-3.77). Consistent with prior work (e.g., Dreber and Johannesson, 2008; Friesen and Gangadharan, 2012), we also find a significant negative effect of the male gender on truthfulness; in the most complete (right-most) models, the gender-deception interaction term is statistically significant (z=2.19 and z=2.27) and indicates that males have an estimated 39 to 46 percent lower excess propensity for truthfulness relative to females in our sample.

6 Conclusion

Vanberg (2017) argues that, due to the strategic incentives first identified by Sutter (2009), Erat and Gneezy's (2012) landmark "white lie" experiments do not provide evidence for pure lie aversion, and indeed, that a new experiment which could reveal pure lie aversion provides no evidence for any lie aversion whatsoever. His general conclusion is that subjects are pragmatic and willing to lie when "it is obvious the (moral) rule does not serve any purpose" (p.458).

Our experiments suggest two revisions to these conclusions: (1) the new Vanberg experi-

[28] As in the white lie experiment, a high proportion of our "black lie" participants predict substantial rates at which Receivers follow their messages (Table 4).

ment does not demonstrate an absence of lie aversion per se, but rather the predominance of another moral imperative, the promotion of collective (Team) interests; and (2) controlling for strategic incentives, but withholding the added Team ethic at play in Vanberg's new experiment, the original conclusions of EG are supported. That is, evidence for pure lie aversion withstands the strategic incentive critique.

Overall, the experiments indicate that subjects weigh competing moral and monetary considerations when making their decisions. A key element of this decision-making, when deception is involved, is a pervasive pure aversion to lies among a large proportion of experimental subjects.

7 References

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

Battigalli, P. & M. Duwfenberg. 2007. "Guilt in Games." *American Economic Review*, 97(2), 170-176.

Battigalli, P., G. Charness, & M. Duwfenberg. 2013. "Deception: The Role of Guilt." *Journal of Economic Behavior & Organization*, 93, 227-232.

Carr, A.. 1968. "Is Business Bluffing Ethical?" *Harvard Business Rev.*, Jan.-Feb., 2-8.

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

Dreber, A. & M. Johannesson. 2008. "Gender Differences in Deception." *Economics Letters*, 99, 197-199.

Edger, M. 2010. "Building Team Cohesion for Success." *Sports Psychology Today*, Nov. 4.

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

Erat, S. & U. Gneezy. 2012. "White Lies." *Management Science*, 58 (4), 723-733.

Festinger, L. 1957. *A Theory of Cognitive Dissonance.* Stanford, CA: Stanford U. Press.

Fischbacher, U. & F. Föllmi-Heusi. 2013. "Lies in Disguise - An Experimental Study on Cheating." *Journal of the European Economic Association*, 11(3), 525-547.

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

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

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

Gneezy, U., A. Kajackaite, & J. Sobel. 2016. "Lying Aversion and the Size of the Lie." Working Paper, U.C. San Diego and U.S.C..

Gneezy, U., B. Rockenbach, & M. Serra-Garcia. 2013. "Measuring Lie Aversion." *Journal of Economic Behavior & Organization*, 93, 293-300.

Hackman, R. 2004. "What Makes for a Great Team?" *American Psych. Association*, June.

Hurkens, S. & N. Kartik, 2009. "Would I Lie to You? On Social Preferences and Lying Aversion." *Experimental Economics*, 12(2), 180-192.

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

Kant, I. 1785. "Fundamental Principles of the Metaphysics of Morals."

Lopez-Perez, R. & E. Spiegelman. 2013. "Why do People Tell the Truth? Experimental Evidence for Pure Lie Aversion." *Experimental Economics*, 16(3), 233-247.

Lundquist, T., T, Ellingsen, E. Gribbe & M. Johannesson. 2009. "The Aversion to Lying." *Journal of Economic Behavior & Organization*, 70 (1-2), 81-92.

Mahon, J. 2016. "The Definition of Lying and Deception." *The Stanford Encyclopedia of Philosophy*, Spring 2016 Edition, Edward N. Zalta, ed..

Mills, J. 1958. "Changes in Moral Attitudes Following Temptation." *J. Personality*, 26(4), 517-31.

Riggio, R. 2011. "10 Rules for High Performing Teams." *Psychology Today*, Aug. 22.

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.

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, 47-60.

Vanberg, C. 2017. "Who Never Tells a Lie?" *Experimental Economics*, 20, 448-459.

Table 1. Vanberg Experiment Results

A. Number of Observations		Truth Disadvantageous (V<W)			
		Full Sample (N=32)		Restricted Sample (N=25)[+]	
		Truth	Untruth	Truth	Untruth
Truth Advantageous (V>W)	Truth	5[++]	24	4[++]	20
	Untruth	0	3	0	1

B. Percentages	Full Sample (N=32)		Restricted Sample (N=25)[+]	
	Percent	z-statistic[+++]	Percent	z-statistic[+++]
Always Truthful Subjects	0.1562	2.433**	0.1600	2.182**
Always Truthful Subjects who Predict Receiver Follow	0.1250	2.138**	0.1200	1.846*

[+]Restricted Sample only includes subjects rating experiment clear/understood.

[++] 4 of 5 truth-tellers predict Receiver follow in the full sample, and 3 of 4 truth-tellers predict Receiver follow in the Restricted Sample.

[+++]z-statistic for null of zero percent.

**, * Significant at 5%, 10% (two-sided).

Table 2. Interpreting an Advantageous Untruthful Message in the Vanberg Experiment

Statement ↓	% Agree (N=31) (1)	% Disagree (N=31) (2)	Difference (1)-(2)	z-statistic[+]
Helping the Team	0.6774	0.1935	0.4839	5.460***
A Lie	0.3871	0.2903	0.0968	0.937
Enabling Teammate to Earn More Money	0.6129	0.1613	0.4516	4.802***
The Right Thing to Do	0.5161	0.2903	0.2258	2.273**

[+]z-statistic for null of zero difference in percentages (% Agree - % Disagree).

***, ** Significant at 1%, 5% (two-sided).

Table 3. White Lie Results

	Percentage	z statistic[++]
Percent Truthful (Full Sample, N=54)	40.74%	6.093***
Percent Truthful (Restricted Sample)[+]	44.44%	6.000***
Percent Male[+++]	62.96%	
Percent Predicting Receiver Follow Rate > 50%[+++]	72.22%	
Percent Predicting Receiver Follow Rate > 25%[+++]	94.44%	
Percent Rating Experiment Clear/Understood[+++]	83.33%	

[+]Restricted Sample only includes subjects rating experiment clear/understood (N=45).

[++]z-statistic for null of zero percent truthful. ***Significant at 1% (two-sided).

[+++]Full Sample (N=54).

Table 4. Black Lie Results

	Deception Game (1)	Dictator Game (2)	Difference (3)
	% (No. of Obs.)	% (No. of Obs.)	(1)-(2) (z-statistic)[++]
Truthful/Generous (Full Sample)	0.800 (55)	0.442 (52)	0.358 (4.089)***
Truthful/Generous (Restricted Sample)[+]	0.783 (46)	0.395 (38)	0.388 (3.881)***
Male	0.509 (55)	0.577 (52)	-0.068 (-0.706)
Percent Predicting Receiver Follow Rate > 50%	0.545 (55)		
Percent Predicting Receiver Follow Rate > 25%	0.891 (55)		
Percent Rating Experiment Clear/Understood	0.836 (55)	0.731 (52)	0.105 (1.333)

[+]Restricted Sample includes subjects rating experiment clear/understood. [++]z-statistic for null of zero difference in percentages (Deception Game – Dictator Game). ***Significant at 1% (two-tail).

Table 5. Black Lie Regressions

	OLS	Probit	OLS	Probit
	Coefficient (z-statistic)	Marginal Effect (z-statistic)	Coefficient (z-statistic)	Marginal Effect (z-statistic)
Deception Dummy	0.356 (3.930)***	0.356 (3.766)***	0.573 (4.724)***	0.577 (4.075)***
Male Dummy	0.011 (0.126)	0.003 (0.036)	0.220 (1.532)	0.214 (1.548)
Male*Deception	No	No	-0.395 (-2.192)**	-0.463 (-2.267)**
Course Effects	Yes	Yes	Yes	Yes

Dependent Variable: Truthful / Generous (0-1). N=107. Robust z-statistics in parentheses. **,*** Significant at 5%, 1% (two-tail).

Figure 1. Vanberg Experiment: Percent Truthful/Untruthful When Truth (Untruth) Is Advantageous

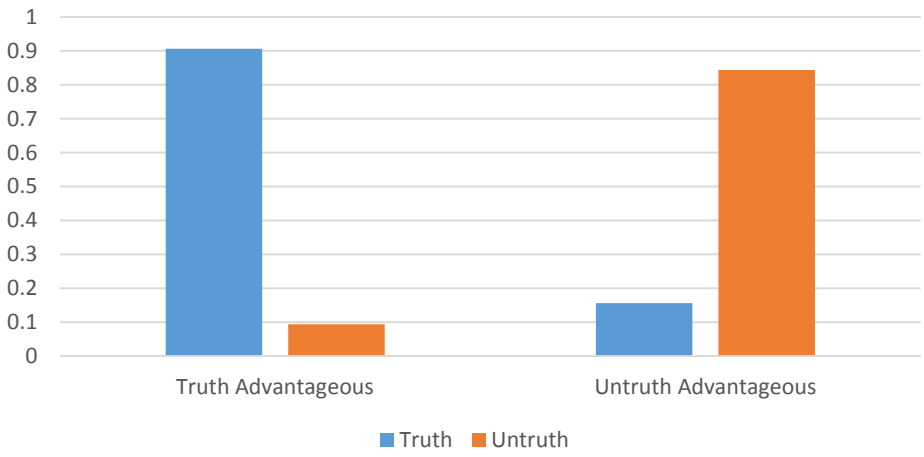

Figure 2. Vanberg Experiment: Interpretation of Advantageous Untruthful Message

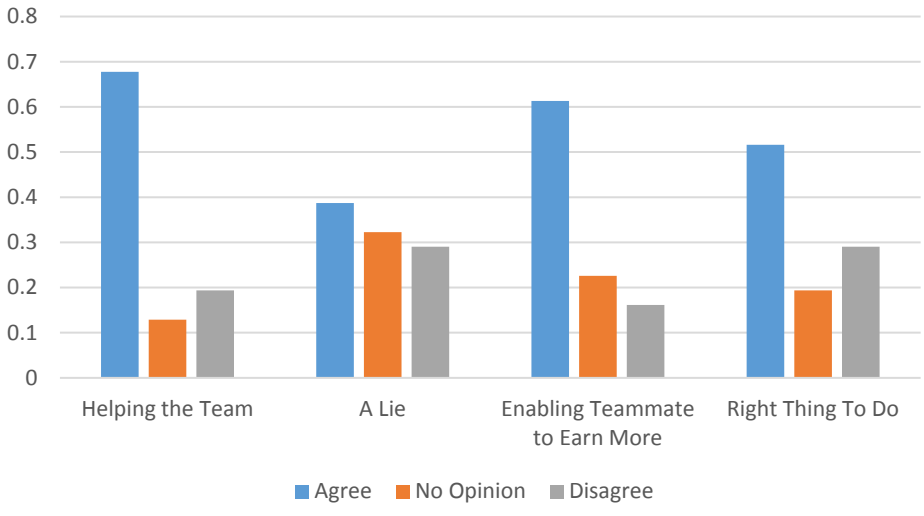

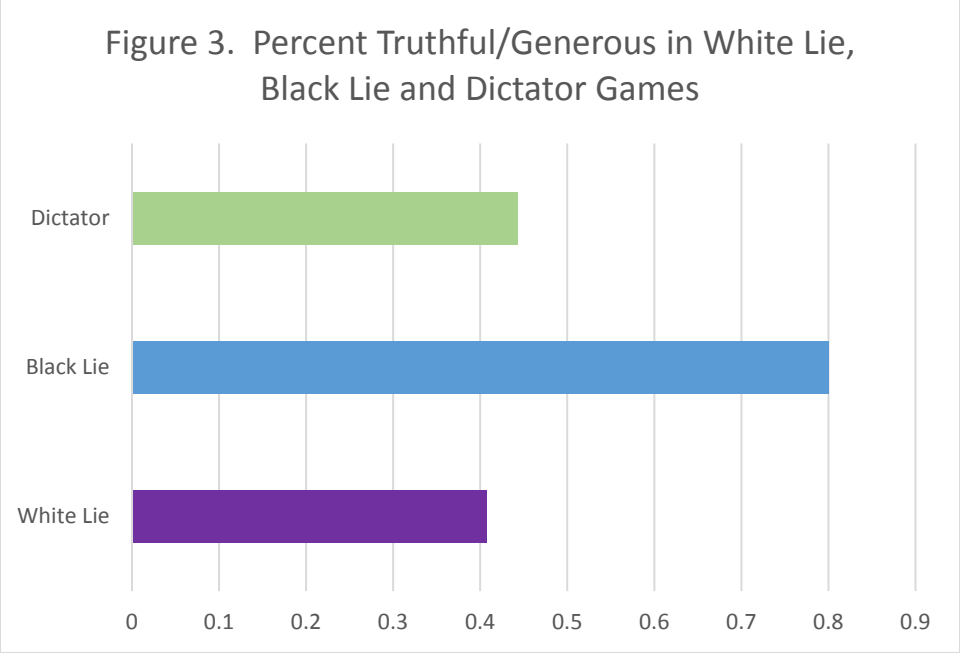
Figure 3. Percent Truthful/Generous in White Lie, Black Lie and Dictator Games
Dictator
Black Lie
White Lie
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9