

Deception: The Role of Consequences

By URI GNEEZY*

Deception is part of many economic interactions. Business people, politicians, diplomats, lawyers, and students in the experimental laboratory who make use of private information do not always do so honestly. This observation indicates that behavior often rejects the moral approach to deception. As St. Augustine wrote, "To me, however, it seems certain that every lie is a sin. . ." (St. Augustine, 421). Later, philosophers like Immanuel Kant (1787) again adopted this uncompromising moral stance when arguing against lying.

At the other extreme, economic theory is built on the assumption of "homo economicus," a figure who acts selfishly and is unconcerned about the well-being of others.¹ An implication of this assumption is that lies will be told whenever it is beneficial for the liar, regardless of their effect on the other party.² Another implication is that there is no negative outcome as-

sociated with lying per se. This assumption is very useful in many economic models. Consider contract theory, where it is assumed that without an explicit contract, neither side will fulfill its respective obligations. For example, George Akerlof's (1970) paper on asymmetric information and the market for lemons assumes that sellers of used cars will always lie if it is in their benefit to do so.

In the mechanism design literature (e.g., Bengt Holmstrom, 1979), the standard assumption is that people will tell the truth only if this is incentive-compatible given material outcomes. In the literature on tax evasion, the choice of whether to avoid paying taxes is considered a decision under uncertainty; cost is treated as a product of the probability of being caught and the cost of punishment, whereas benefit is simply the money saved by avoiding payment. However, there is no cost associated with the very act of lying (Michael Alingham and Agnar Sandmo, 1972). Another example is the game theoretic treatment of "cheap talk" (Crawford and Joel Sobel, 1982).

An intermediate approach is taken by utilitarian philosophers (e.g., Jeremy Bentham, 1789). Utilitarianism prescribes that, when choosing whether to lie, one should weigh benefits against harm, and happiness against unhappiness. As Martin Luther stated, "What harm would it do, if a man told a good strong lie for the sake of the good and for the Christian church. . . a lie out of necessity, a useful lie, a helpful lie, such lies would not be against God, he would accept them."³ Similarly to the economic theory approach, this type of calculation implies that lies, apart from their resultant harm and benefit, are in themselves neutral. A lie and a truthful statement that achieve the same monetary payoffs (for both sides) are considered

* Graduate School of Business, University of Chicago, 5807 South Woodlawn Avenue, Chicago, IL 60637 (e-mail: uri.gneezy.gsb.uchicago.edu). I thank Douglas Bernheim and two anonymous reviewers for insightful comments that considerably improved the paper. I also thank Andreas Blume, Gary Charness, Rachel Croson, Martin Dufwenberg, Georg Kirchsteiger, David Levine, Muriel Niederle, Yuval Rottenstreich, Maurice Schweitzer, Richard Thaler, George Wu, and seminar participants at numerous universities for their comments and suggestions. Ira Leybman provided valuable help in running the experiment. I became interested in deception when my father was terminally ill and his physicians created in him a belief that they considered to be untrue. I dedicate this paper to his memory.

¹ Important deviations from this assumption in economic modeling are found in Kenneth Arrow's (1972) discussion of trust, Gary Becker's (1976) modeling of altruistic preferences, and Akerlof's (1982) study of the fair-wage hypothesis. For a general discussion, see Becker (1993): "The economic approach I refer to does not assume that individuals are motivated solely by selfishness or material gain. It is a *method* of analysis, not an assumption about particular motivations. Along with others, I have tried to pry economists away from narrow assumptions about self-interest. Behavior is driven by a much richer set of values and preferences" (p. 385).

² Note that this does not mean that a completely selfish person will always lie. There may be strategic reasons not to lie. For example, see the David Kreps and Robert Wilson

(1982) discussion of reputation and imperfect information; see also Vincent P. Crawford (2003).

³ Cited by his secretary, in a letter in Max Lenz, ed., *Briefwechsel Landgraf Phillips des Grossmuthigen von Hessen mit Bucer*, Vol. 1.

equivalent (Sissela Bok, 1978, Ch. 4). This is a consequentialist approach. An alternative approach—one that distinguishes between two decisions with the same payoff set according to the process leading to the outcomes—would be called nonconsequentialist.

St. Augustine's approach is normative in the sense that it prescribes "what a person should do."⁴ His injunction is (unfortunately?) not supported by a casual observation of real life: people do lie. Economic theory is normative in the sense that it prescribes "what a rational economic agent should do." This approach is also not supported by casual observation: even economists tell the truth from time to time, in the absence of any strategic justification for doing so. The utilitarian approach predicts that if people do care about the well-being of others, the decision to lie (or not) may depend on its cost to the other side. As a result, people will not go to either extreme of always lying or always telling the truth. As I show below, however, people do distinguish between lying and "innocent" choices, even when the decisions do not differ in monetary outcomes. In particular, people are less likely to choose the outcome that maximizes their own monetary payoff if it involves a lie than if it involves an innocent choice. Hence, the consequentialist assumption of utilitarianism is rejected.

I empirically studied the role of consequences in the decision concerning whether to lie.⁵ I considered a two-person interaction in

which lying increases the payoffs to the liar at the expense of her counterpart, and I asked the following question: How do changes in relative payoffs influence this decision? The main empirical finding is that people not only care about their own gain from lying; they also are sensitive to the harm that lying may cause the other side. The average person prefers not to lie, when doing so only increases her payoff a little but reduces the other's payoff a great deal.

I. Classification of Lies, and a Definition

It is interesting to note that the literature offers many ways to classify lies.⁶ I base my classification on the consequences that the lie produces. Using this criterion, one can devise four major categories. First, there are lies that help both sides, or at least do not harm anyone, for instance, a white lie that costs the liar nothing and makes the counterpart feel good ("You look great today!"). In the second category I place lies that help the other person even if it harms the liar. The motivation for this kind of lie may be pure altruism (Becker, 1976), an impure motive according to which people enjoy the act of giving (James Andreoni, 1990), or an "efficiency motive," according to which people prefer outcomes that enlarge total surplus (Gary Charness and Matthew Rabin, 2002). In the third category are lies that do not help the liar but can harm both sides or, at the very least, the other person. The motive for this might be a spiteful reaction to unfair behavior.

The fourth category I consider in this paper includes lies that increase the payoff to the liar and decrease the payoff to the other party. I argue that this is the relevant category for many economic events, such as those covered by mechanism design and contract theory. If person A signs a contract with person B, it is simply to prevent B from acting in ways that

⁴ Although Augustine was categorically against lies, he distinguished between different types of lies. The continuation of the citation reads, "... though it makes a great difference with what intention and on what subject one lies." Similarly, although Jewish texts prohibit lying, certain lies, especially those told to preserve household unity, are regarded as exceptions (Lewis Jacobs, 1960).

⁵ Many other aspects of deception are studied in the literature. Psychologists study personality characteristics of honesty (Hugh Hartshorne and Mark May, 1928), how to detect lies (Paul Ekman, 1992; Albert Vrij, 2001), etc. See Charles Ford (1995) for an introduction to the psychology of deceit, Roy Lewicki (1983) for a behavioral model, and Bella DePaulo et al. (1996) for taxonomy of lies and their classifications according to content, motivation, and magnitude. In accounting, John Evans et al. (2001) examine how preferences for wealth and honesty affect managerial reporting (see also the discussions by Stanley Baiman and Barry Lewis, 1989; Kenneth Koford and Mark Penno, 1992). In experimental economics, Andreas Ortmann and Ralph

Hertwig (2002) compare the costs and benefits of deceiving participants in laboratory experiments, concluding that experimental economists' prohibition of deception is sensible. In business, research focuses on deception in negotiations (e.g., Maurice Schweitzer and Rachel Croson, 1999). See Bok (1978, particularly chapter 4) for what I find the most thoughtful modern treatment on the morality of deception.

⁶ See Lewicki (1983), DePaulo et al. (1996), and Schweitzer (2001).

will increase her payoff at the expense of A's payoff. The focus in this paper is on this fourth category, since I am interested in the cost-benefit relationship between the deceiver's gain and the other's loss. The discussion will be devoted to the influence of anticipated consequences resulting from this kind of lying on the decision whether or not to lie. Hence, the working definition of deception I use (based on Vrij, 2001) is:

"A successful or unsuccessful deliberate attempt, without forewarning, to create in another a belief that the communicator considers to be untrue in order to increase the communicator's payoff at the expense of the other side."

II. Empirical Findings

In order to study empirically the effect of consequences on behavior, I conducted experiments in which I altered the absolute and relative consequences of lies and then measured the effect of this change on participants' propensity to lie. In addition, finding out what people say about the importance of costs to the receiving end is interesting. To study attitudes toward fairness in lying, I constructed questionnaires involving "real life" scenarios. This method is well established in psychology and has been used in economics (e.g., the study of fairness in Daniel Kahneman et al.).

A. A Cheap Talk Sender-Receiver Game

The Cheap Talk Game.—I use a two-player communication game in which one player has private information and the other takes an action. The message precedes the action. Payoffs to both players depend on the action chosen, not on the message. This type of situation can be modeled using a cheap talk sender-receiver game. In these games, communication is the crucial link between the private information and the action. The theoretical question is what form of communication is expected, and how it should affect the action (Crawford and Sobel, 1982).

In the game I studied, there were two possible monetary distributions: A or B. Only player 1 was informed about the monetary consequences

of each option, and the rules of the game were made known to both participants. Player 1 sent one of two possible messages to player 2:

Message A: "Option A will earn you more money than option B."

Message B: "Option B will earn you more money than option A."

After receiving one of these messages, player 2 chose the option to be implemented.

With standard preferences and conflicting objectives (and players who know this), the only cheap talk equilibrium is one in which the message contains no relevant information for the receiver. With a rich distribution of possible payoffs, it is possible to sustain a signaling equilibrium where the advice is followed. (Consider the extreme where the incentives of the two players are fully aligned). Andreas Blume et al. (2001) study the relationship between the incentives' alignment and the form of communication, showing an important correlation between the two. In the game studied in the current paper, the receiver is given no information regarding the alignment of incentives.⁷

I am primarily interested in the message sent by the sender, and hence in the sender's beliefs regarding the effect of the message on player 2's choice of action. To learn this, 50 participants assigned to the role of sender were asked to guess how the receiver would react to their message (they were paid for accuracy). Of these, 41 (82 percent) said that they expected the receiver to follow their message. This indicates some heterogeneity of beliefs among senders, but also that the majority expected a mechanical response by the receiver. As it turned out, 78 percent of the participants who were assigned to the role of receiver followed the sender's message and chose the option that the sender told them would earn them more money. That is, the receiver chose the option "recommended" by the sender in most cases.

For the purpose of the discussion that follows, I analyze the game as a decision problem

⁷ Note that any technical description of the game should specify the receiver's beliefs about the alignment of incentives (see, e.g., Crawford, 2003).

for the sender in the following sense: I assume that (most) senders expect the receiver to be credulous. That is, senders correctly anticipate a mechanical response by receivers. If that is the case, and the sender is selfish, then she will always send the message recommending the outcome that maximizes her own expected payoff.

To further test this assumption, the treatment was repeated with another group of 50 participants assigned to the role of sender. After making their choices, they were told that we had already conducted the experiment with the receiver, and that the receiver they were matched with had chosen to follow the message they had sent.⁸ They were then asked whether they wished to reconsider their previous choice. Three (6 percent) chose to change their message. One moved from telling the truth to lying, and two moved the other way.

To conclude, within the context of the experiment, if the sender is interested simply in maximizing her own payoff, and she has rational expectations about the reaction of the receiver to the message she sends, she should always lie. Moreover, the sender understands this. This result is in line with Blume et al. (2001), who found that even when equilibrium prescribes that the receiver ignore the message, she tends to follow it (what they call “sucker behavior”). I find this property instructive because it helps separate strategic motives from fairness motives. Because the sender expects the lie to “work,” her only concern relates to the fairness of lying.

Procedure.—The participants were 450 undergraduate students at Technion University and the University of Haifa who volunteered to participate in the experiment after class. They were told that the experiment would take about 15 minutes, and that everyone would be paid. In the instructions (see Appendix) for both player 1 and player 2, it was written that there were two possible outcomes to the experiment. Although the actual choice between the options was to be made by player 2, only player 1 was informed about the monetary consequences of

TABLE 1—THE DIFFERENT PAYOFFS USED IN THE DECEPTION GAME

Treatment	Option	Payoff to	
		Player 1	Player 2
1	A	5	6
	B	6	5
2	A	5	15
	B	6	5
3	A	5	15
	B	15	5

each option. The rules were made known to both participants. The participants were told that neither of them would ever know the identity of the other. Identification with respect to the experimenter was established using students’ ID numbers.

After reading the instructions, player 1 was asked to send one of two possible messages to player 2, as described above. This message was the only information player 2 had about the payoffs.⁹ In all three treatments, payoffs were constructed such that if option A were chosen, player 1 earned less than if option B were chosen and the reverse for player 2 (recall that player 2 did not know that the payoffs were inverse). As a result, message B was not true, and the payoff associated with it for player 1 was larger than the payoffs associated with message A. Actual payoffs used in the experiment are presented in Table 1 (75 pairs of participants per cell).

Results.—The results of the experiment, in terms of the fraction of player 1s who lied, are presented in Figure 1. The figure is constructed according to sender’s profit and receiver’s loss.

In treatment 1, where the gain for player 1, from lying, was \$1 and the loss for player 2 was also \$1, 27 (36 percent) of the 75 senders lied. In treatment 2, where senders still gained \$1 from lying, but the loss for player 2 was increased to \$10, the number of participants who lied declined to 13 (17 percent). Finally, in treatment 3, where profit to sender and cost to

⁸ The original instructions were adapted such that this would not contradict what they had been told previously and to prevent deception by the experimenter.

⁹ Different types of message delivery can affect the outcomes. For example, Kathleen Valley et al. (1998) study bilateral bargaining with asymmetric information and find different degrees of truth-telling and trust across different mediums of communication.




FIGURE 1. FRACTION OF PARTICIPANTS WHO LIED IN THE DECEPTION GAME

Note: The horizontal axis represents the gains from lying for player 1 and the associated loss for player 2.

receiver were both \$10, that number rose to 39 (52 percent). A statistical comparison of these differences shows that they are all significant.¹⁰

Deception versus Choices between Allocations.—In order to determine the extent to which these results reflect an aversion to lying as opposed to preferences over distributions of payoffs, I used a control dictator treatment in which player 1 chose between two monetary allocations, just as in the deception game. Player 2 has no choice in this control treatment. The probability of executing player 1's choice was 80 percent, while in the other 20 percent the alternative allocation was implemented. Since approximately 80 percent of player 2s followed player 1s' recommendation in the deception game, this results in a treatment that is equivalent, in payoff, to the deception game. If player 1s had chosen the materially advantageous allocation more often in this control treatment, it

would be direct evidence of lie-aversion (and against consequentialist preferences). The results are presented in Figure 2, with $N = 50$ in each of the cells of the dictator game.

The results presented in Figure 2 display the same pattern we observed in the deception game, but to a much greater degree. The results of the two games are compared in Table 2.

From these results I conclude that it is not only care for others that motivates behavior, but also aversion to lying. People's choices reflect nonconsequentialist preferences since, for example, they treat the choice between (5, 6) and (6, 5) differently, depending on whether it was a simple choice or a lie that led to the final outcome.

B. The Questionnaires

What do people think about the role of consequences in lying, and what do they say about the relative fairness of different lies? I studied these issues with a set of questionnaires whose items referred to an empirically realistic scenario. The participants in this study were students at the University of Chicago who volunteered to fill out the questionnaires and were paid \$1 for their participation. They were asked to judge the following scenario:

¹⁰ The p -values are approximated to three decimal places and calculated from a one-tailed test of the equality of proportions, using normal approximation to the binomial distribution. For the comparison of treatment 1 and 2, $Z = 2.58$, and $p = .005$. For treatment 1 versus 3, $Z = 1.97$, and $p = .024$, and for treatment 2 versus treatment 3, $Z = 4.48$ and $p = .001$.

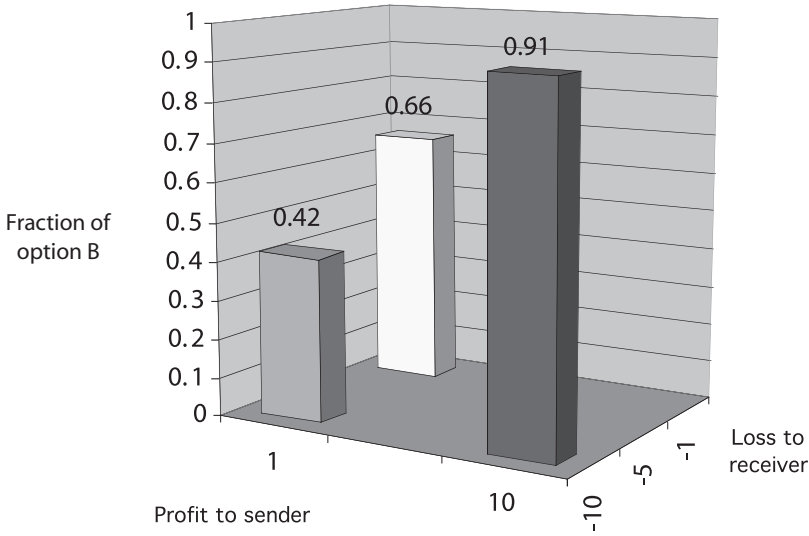


FIGURE 2. FRACTION OF PLAYER 1S WHO CHOSE OPTION B IN THE DICTATOR GAME

Note: The horizontal axis represents the gains from choosing B for player 1 and the associated loss for player 2.

TABLE 2—THE FRACTION OF PLAYER 1S WHO CHOSE ALLOCATION B

Game	Allocations		
	5, 6 versus 6, 5	5, 15 versus 6, 5	5, 15 versus 15, 5
Deception	0.36	0.17	0.52
Dictator	0.66	0.42	0.90

Notes: All differences between the dictator game and the deception game for a given distribution of payoffs are statistically significant at $P < 0.01$. Differences between the different allocations within the dictator game are also statistically significant at the 0.01 level.

Mr. Johnson is about to close a deal and sell his car for \$1,200. The engine’s oil-pump does not work well, and Mr. Johnson knows that if the buyer learns about this, he will have to reduce the price by \$250 (the cost of fixing the pump). If Mr. Johnson doesn’t tell the buyer, the engine will overheat on the first hot day, resulting in damages of \$250 for the buyer. Being winter, the only way the buyer can learn about this now is if Mr. Johnson were to tell him. Otherwise, the buyer will learn about it only on the next hot day. Mr. Johnson chose not to tell the buyer about the problems with the oil pump. In your opinion, Mr. Johnson’s behavior is

(please circle one): completely fair; fair; unfair; very unfair.

What would your answer be if the cost of fixing the damage for the buyer in case Mr. Johnson does not tell him is \$1,000 instead of \$250? Mr. Johnson’s behavior is (please circle one): completely fair; fair; unfair; very unfair.

Although there was no difference between the two scenarios in terms of the seller’s pay-offs, the buyer’s cost increases from \$250 to \$1,000. I used both a between-subjects design (i.e., “what would be. . .”), with $N = 50$ students answering each question, and a within-subjects design (i.e., the participants answered the question for both parameters as they are presented above); again $N = 50$. The students’ responses are presented in Figure 3.

The difference between the answers to the first and second question in the between-subjects design is significant ($p < .05$).¹¹ Inspection of the within-subjects design shows a large difference in choices. In the \$250 cost question, 70 percent of the participants chose “unfair” and 18 percent chose “very unfair.” In

¹¹ Using both Kolmogorov-Smirnov and Wilcoxon rank-sum tests.

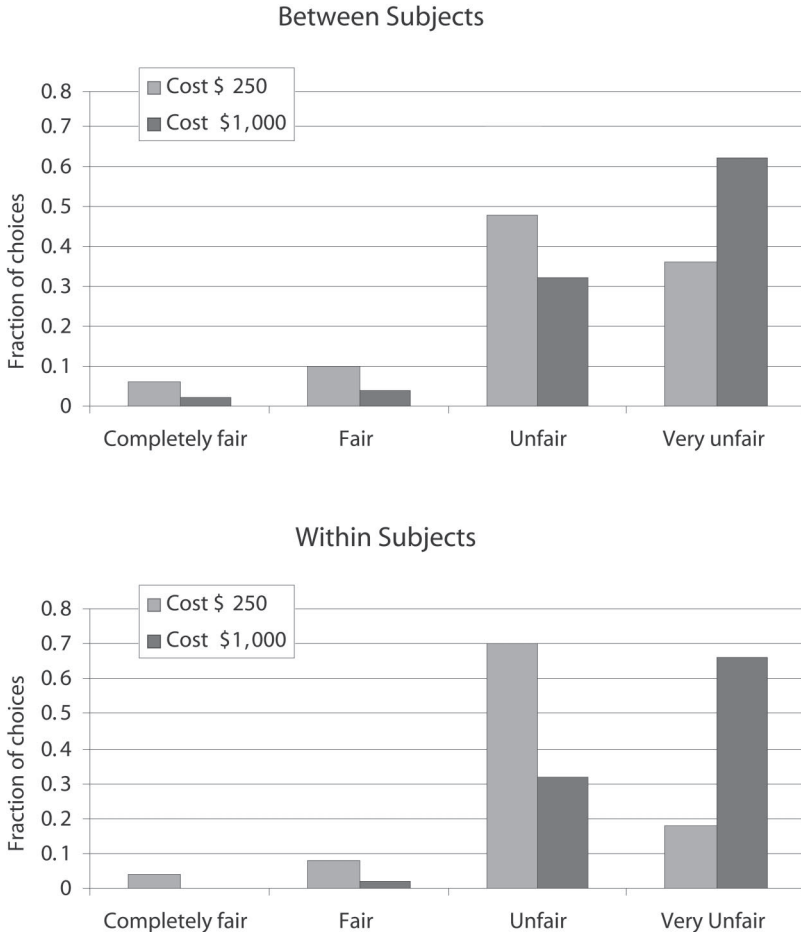


FIGURE 3. REPLIES TO THE CAR SALE QUESTION

Note: Purchase price is \$1,200, seller cost of repair is \$250, and buyer cost of repair is either \$250 or \$1,000.

the \$1,000 question, only 32 percent chose “unfair,” but 66 percent chose “very unfair.” This difference is highly significant ($p < 0.001$).¹²

The basic finding here is that people think it is less fair to lie as the cost for the other side

¹² The differences between the replies to the first and second question are smaller in the between-subjects design. In the within-subjects design, people apparently wanted to emphasize that lying is worse when the costs for the buyer were higher, and consequently they answered only “unfair” to the first question. We see that in the within-subjects design 70 percent chose “unfair” to the first questions, while in the between-subjects design only 48 percent chose this response. The difference between the replies in the within-subjects design and between-subjects design is statistically significant at 0.05.

increases: 30 (60 percent) out of the 50 participants indicated that the lie was less fair (more unfair) when the cost was higher. This intuition is strengthened when we consider a third scenario in which the problem lies with the brakes, and the risk to the buyer is an automobile accident. When this question was asked of 20 subjects at the University of Chicago, they found it insulting and unworthy of an answer. These results were replicated with other scenarios, such as questions regarding the purchase of a house.¹³

¹³ The results, as well as the entire data set for the above question, can be found at: <http://gsbwww.uchicago.edu/fac/uri.gneezy/vita/>.

III. Discussion

When do people lie? This paper sheds some light on the propensity of people to lie in situations where no penalty is associated with lying. The analysis is based on consequences, that is, changes in wealth resulting from a lie. These consequences turn out to have an important effect on behavior. The first result is that people are sensitive to their gain when deciding to lie. Second, people care not only how much they gain from a lie, but also how much the other side loses. This unselfish motive diminishes with the size of the gains to the decision maker herself.

The implications of these results are illustrated by the purchase of a car: you can trust what the seller says about the condition of the brakes more than what she says about the state of the air conditioning. This result may also explain why people are more accepting of fraudulent behavior directed at large organizations or rich counterparts than at individuals: the monetary cost may be identical, but the damage to the individual is perceived as greater. For example, people are more accepting of lies told by an employee to an employer than vice versa (David Strahlberg, 2001) and are more likely to deceive insurance companies than private citizens (Sharon Tennyson, 1997; Insurance Research Council, 1991).

Extending the standard model using empirical evidence can help us understand important economic phenomena. Some attempts to model deception are inconsistent with the results presented here. First, assuming that people are either completely honest or not at all is problematic. For example, in Kenneth Koford and Mark Penno (1992), agents are one of two types: "ethical" (fully honest) or "economic" (willing to tell any lie necessary to maximize wealth). Ethical types never lie because they experience infinite disutility from lying, whereas economic types always lie to maximize their wealth because they experience no disutility from lying. This model cannot explain why people are sensitive to payoffs associated with unethical behavior. Second, models that assume the decision maker computes a simple cost-benefit analysis of her own monetary payoffs prior to deciding whether to lie are also inconsistent with the results. For example, Stanley

Baiman and Barry Lewis's (1989) threshold model assumes that individuals experience a small fixed disutility from lying, i.e., that they are honest for all payoffs below their personal disutility threshold and lie to maximize wealth for all payoffs at or above the threshold. This model cannot explain the difference in behavior observed when the cost for the decision maker is fixed.

A third type of modeling can be based on formal models of social preferences that assume people are self-interested but also concerned about the payoffs to others. Distributional models in which an agent's preferences are influenced by the final distribution of payoffs are presented in Georg Kirchsteiger (1992), Ernst Fehr and Klaus Schmidt (1999), and Gary Bolton and Axel Ockenfels (2000). It is assumed in these models that, for a given own-material payoff, a person's utility decreases with the difference between the own-payoff and that of the counterpart. When one tries to use a distributional model for deception games, however, it generates some strong unintuitive predictions. For example, consider the case where player 1 has to choose between lying to player 2, thereby obtaining the following payoffs (\$6 for herself and \$5 for the other), and telling the truth to obtain the following payoffs (\$5 for herself and \$15 for the other). According to the Fehr and Schmidt (1999) model, letting x be the decision maker's payoff and y the other person's payoff, the decision maker's utility U is $x - \alpha_{\max}\{0, y - x\} - \beta_{\max}\{0, x - y\}$, where α and β are restricted such that $0 \leq \beta < 1$, $\beta \leq \alpha$. We get: $U(6, 5) = 6 - \beta$, $U(5, 15) = 5 - 10\alpha$, and $U(6, 5) - U(5, 15) = 1 - \beta + 10\alpha$. This last term is always positive given the restrictions above, meaning that the model predicts the choice of lying resulting in (6, 5) over truth telling resulting in (5, 15). Moreover, keeping all else constant, the more player 2 earns in the second option, the smaller the probability that player 1 will choose it. This prediction is rejected by the results presented above.¹⁴

A more promising direction for modeling deception would take into account the fact that

¹⁴ The purely distributional aspect of the Charness and Rabin (2002) model would predict the qualitative results that people are averse to reducing the total surplus.

people are not indifferent to the process leading up to the outcome. Decision makers apply specific criteria when deciding whether to lie, similar to the letdown aversion findings in Martin Dufwenberg and Gneezy (2000) and the guilt aversion findings in Gary Charness and Dufwenberg (2003). That is, people have nonconsequential preferences in which they treat the same monetary outcome differently, depending on the process that leads up to it. Moreover, different people weigh these preferences differently. In line with this, focusing on the person being deceived, Jordi Brandts and Charness (2003) found that the willingness to punish an unfair action is sensitive to whether this action was preceded by a deceptive message. Under such intrinsic motivation to tell the truth, penalizing deception may be counterproductive and result in even more lies (Gneezy and Aldo Rustichini, 2000).

This could explain, for example, one empirical departure from the result of Akerlof's (1970) market-for-lemons model. One of the conclusions raised by the lemons model is that only dealers who can offer a warranty will sell used cars of high quality. This conclusion is not in line with the real-world coexistence of professional car dealers and private sellers who sell cars of high quality without a warranty. An explanation based on the findings of the current paper may suggest that car dealers have lower costs associated with lying than private sellers. People trust the honesty of private sellers more than they trust professional car dealers.¹⁵ As a result, they are willing to buy a used car without a warranty from the former but not from the latter.

APPENDIX A: INSTRUCTIONS FOR EXPERIMENT, PLAYER 1

This is a short experiment in decision making. In this experiment, you will be matched with another student from a different class. Neither of you will ever know the identity of the other. The money that you earn will be paid to you next week, privately and in cash.

¹⁵ When asked, 85 percent of 40 MBA students stated this belief.

Two possible monetary payments are available to you and your counterpart in the experiment. The two payment options are:

Option A: \$5 to you and \$6 to the other student

Option B: \$6 to you and \$5 to the other student

The choice rests with the other student who will have to choose either option A or option B. The only information your counterpart will have is information sent by you in a message. That is, he or she will not know the monetary payments associated with each choice.

We now ask you to choose one of the following two possible messages, which you will send to your counterpart:

Message 1: "Option A will earn you more money than option B."

Message 2: "Option B will earn you more money than option A."

We will show the other student your message, and ask him or her to choose either A or B. To repeat, your counterpart's choice will determine the payments in the experiment. However, your counterpart will never know what sums were actually offered in the option not chosen (that is, he or she will never know whether your message was true or not). Moreover, he or she will never know the sums to be paid to you according to the different options.

We will pay the two of you according to the choice made by your counterpart.

I choose to send (please circle one option):

Message 1 Message 2

APPENDIX B: INSTRUCTIONS FOR EXPERIMENT, PLAYER 2

This is a short experiment in decision-making. In this experiment you will be matched with another student from a different class. Neither of you will ever know the identity of the

other. The money that you earn will be paid to you next week, privately and in cash.

Two possible monetary payments are available to you and your counterpart in the experiment. The payments depend on the option chosen by you. We showed the two payment options to your counterpart. The only information you will have is the message your counterpart sends to you.

Two possible messages could be sent:

Message 1: "Option A will earn you more money than option B."

Message 2: "Option B will earn you more money than option A."

Your counterpart decided to send you message: _____

We now ask you to choose either option A or option B. Your choice will determine the payments in the experiment. You will never know what sums were actually offered in the option not chosen (that is, if the message sent by your counterpart was true or not). Moreover, you will never know the sums your counterpart could be paid with the other option.

We will pay the two of you according to the choice you make.

I choose (please circle one):

Option A Option B

REFERENCES

- Akerlof, George A.** "The Market for "Lemons": Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 1970, 84(3), pp. 488–500.
- Akerlof, George A.** "Labor Contracts as a Partial Gift Exchange." *Quarterly Journal of Economics*, 1982, 97(4), pp. 543–69.
- Alingham, Michael G. and Agnar, Sandmo.** "Income Tax Evasion: A Theoretical Analysis." *Journal of Public Economics*, 1972, (1) pp. 323–38.
- Andreoni, James.** "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving?" *Economic Journal*, 1990, 100(401), pp. 464–77.
- Arrow, Kenneth J.** "Gifts and Exchanges." *Philosophy and Public Affairs*, 1972, 1(4), pp. 343–62.
- Augustine.** *Enchiridion on faith, hope, and love*. Washington, DC: Regnery Publishing, Inc, 1996.
- Baiman, Stanley and Lewis, Barry L.** "An Experiment Testing the Behavioral Equivalence of Strategically Equivalent Employment Contracts." *Journal of Accounting Research*, 1989, 27(1), pp. 1–20.
- Becker, Gary S.** "Altruism, Egoism, and Genetic Fitness: Economics and Sociobiology." *Journal of Economic Literature*, 1976, 14(3), pp. 817–26.
- Becker, Gary S.** "Nobel Lecture: The Economic Way of Looking at Behavior." *Journal of Political Economy*, 1993, 101(3), pp. 385–409.
- Bentham, Jeremy.** *An introduction to the principles of morals and legislation*. 1789. Oxford: Clarendon Press, 1907.
- Blume, Andreas; DeJong, Douglas V.; Kim, Yong Gwan and Sprinkle, Geoffrey B.** "Evolution of Communication with Partial Common Interest." *Games and Economic Behavior*, 2001, 37(1), pp. 79–120.
- Bok, Sissela.** *Lying: Moral choices in public and private life*. New York: Vintage Books, 1978.
- Bolton, Gary E. and Ockenfels, Axel.** "ERC: A Theory of Equity, Reciprocity, and Competition." *American Economic Review*, 2000, 90(1), pp. 166–93.
- Brandts, Jordi and Charness, Gary.** "Truth or Consequences: An Experiment." *Management Science*, 2003, 49(1), pp. 116–30.
- Charness, Gary and Dufwenberg, Martin.** "Promises and Partnership." University of California, Santa Barbara, 2003.
- Charness, Gary and Rabin, Matthew.** "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics*, 2002, 117(3), pp. 817–69.
- Crawford, Vincent P.** "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions." *American Economic Review*, 2003, 93(1), pp. 133–49.
- Crawford, Vincent P. and Sobel, Joel.** "Strategic Information Transmission." *Econometrica*, 1982, 50(6), pp. 1431–51.
- DePaulo, Bella M.; Kashy, Deborah A.; Kirken-dol, Susan; Wyer, Melissa and Epstein,**

- Jennifer A.** "Lying in Everyday Life." *Journal of Personality and Social Psychology*, 1996, 70(5), pp. 979–95.
- Dufwenberg, Martin and Gneezy, Uri.** "Measuring Beliefs in an Experimental Lost Wallet Game." *Games and Economic Behavior*, 2000, 30(2), pp. 163–82.
- Ekman, Paul.** *Telling lies: Clues to deceit in the marketplace, politics, and marriage*. New York: W. W. Norton & Co, Inc, 1992.
- Evans, John H., III; Hannan, Lynn; Krishnan, Ranjani and Moser, Donald V.** "Honesty in Managerial Reporting." *Accounting Review*, 2001, 76(4), pp. 537–59.
- Fehr, Ernst and Schmidt, Klaus M.** "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics*, 1999, 114(3), pp. 817–68.
- Ford, Charles V.** *Lies! Lies! Lies! The psychology of deceit*. Washington, DC: American Psychiatric Press, Inc, 1995.
- Gneezy, Uri and Rustichini, Aldo.** "A Fine Is a Price." *Journal of Legal Studies*, 2000, 29(1), pp. 1–17.
- Hartshorne, Hugh and May, Mark A.** *Studies in the nature of character, I: Studies in deceit*. New York: Macmillan, 1928.
- Hertwig, Ralph and Ortmann, Andreas.** "Deception in Experiments: The Costs of an Alleged Method of Last Resort." Unpublished Paper, 2004.
- Holmstrom, Bengt.** "Moral Hazard and Observability." *Bell Journal of Economics*, 1979, 10(1), pp. 74–91.
- Insurance Research Council.** "Public Attitude Monitor." Malvern, PA, 1991.
- Jacobs, Lewis.** *Jewish values*. London: Vallentine Mitchell, 1960.
- Kahneman, Daniel; Knetsch, Jack L. and Thaler, Richard.** "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *American Economic Review*, 1986, 76(4), pp. 728–41.
- Kant, Immanuel.** "On a Supposed Right to Lie from Benevolent Motives," 1787, in Lewis W. Beck, ed., *The critique of practical reason and other writings in moral philosophy*. Chicago: University of Chicago Press, 1949, pp. 346–50.
- Kirchsteiger, Georg.** "Interpersonal Comparisons and Behaviour on Markets," Vienna: PhD diss, University of Vienna, 1992.
- Koford, Kenneth and Penno, Mark.** "Accounting, Principal-Agent Theory, and Self-Interested Behavior," in Norman E. Bowie and Edward R. Freeman, eds., *Ethics and agency theory: An introduction*. New York: Oxford University Press, 1992, pp. 127–42.
- Kreps, David and Wilson, Robert.** "Reputation and Imperfect Information." *Journal of Economic Theory*, 1982, 27(2), pp. 253–79.
- Lewicki, Roy J.** "Lying and Deception: A Behavioral Model, with Application to Negotiation," in Max H. Bazerman and Roy J. Lewicki, eds., *Negotiation in organizations*. Beverly Hills: Sage Publication, 1983, pp. 68–90.
- Ortmann, Andreas and Hertwig, Ralph.** "The Costs of Deception: Evidence from Psychology." *Experimental Economics*, 2002, 5(2), pp. 111–31.
- Schweitzer, Maurice E.** "Deception in Negotiation," in Stephen J. Hoch, Howard C. Kunreuther, and Robert E. Gunther, eds., *Wharton on decision making*. New York: John Wiley & Sons, 2003, pp. 187–200.
- Schweitzer, Maurice E. and Croson, Rachel.** "Curtailing Deception: The Impact of Direct Questions on Lies and Omissions." *The International Journal of Conflict Management*, 1999, 10(3), pp. 225–48.
- Strahlberg, David.** "Nice Young Man Goes Wild," *Chicago Business*. November 5, 2001.
- Tennyson, Sharon.** "Economic Institutions and Individual Ethics: A Study of Consumer Attitudes toward Insurance Fraud." *Journal of Economic Behavior and Organization*, 1997, 32(2), pp. 247–65.
- Valley, Kathleen L.; Moag, Joseph and Bazerman, Max H.** "A Matter of Trust: Effects of Communication on the Efficiency and Distribution of Outcomes." *Journal of Economic Behavior and Organization*, 1998, 34(2), pp. 211–38.
- Vrij, Albert.** *Detecting lies and deceit: The psychology of lying and the implications for professional practice*. New York: John Wiley & Sons, 2001.