Deception: The role of consequences

Uri Gneezy*

April 8, 2002

Abstract: This paper studies the role of consequences in a person's decision to lie. Based on findings from an experiment with a deception game, as well as from questionnaires, I propose a simple formulation of preferences to describe deception behavior. The decision maker uses the "truth telling" outcome as a reference level when evaluating the benefits of lying. The monetary consequences of the lie are compared to this reference level. In the formulation used in this paper the decision maker's utility depends on her own intentions. She is selfish in the sense of maximizing her own payoffs, but sensitive to the cost her lie imposes on the other side. Sensitivity diminishes with the size of payoffs. Moreover, since perception of the counterpart's cost is subjective. When there are differences in wealth as in employee-employer relations or a consumer-insurer interactions, the decision maker is more likely to lie the wealthier the counterpart.

* I thank Gary Charness, Martin Dufwenberg, Georg Kirchsteiger, David Levine, Muriel Niederle, Yuval Rottenstreich Maurice Schweitzer, Richard Thaler George Wu, and seminar participants at Berkeley, University of Chicago and Duke for their comments and suggestions. Ira Leybman provided valuable help in running the experiment. The University of Chicago Graduate School of Business and Technion. Email: ugneezy@gsb.uchicago.edu

"*The whole design of free market capitalism depends upon free people acting responsibly. Business people must answer not just to the demands of the market or self-interest, but to the demands of conscience*." [1]

1. Introduction

Deception is part of many economic interactions. Businesspeople, politicians, diplomats, lawyers, and students in the experimental laboratory who make use of private information do not always do so honestly (DePaulo and Kashy, 1998). This observation contradicts 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 Kant (1787) likewise took this extreme moral approach 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.[2] An implication of this assumption is that lies will be told whenever it is beneficial for the liar, regardless of the effect on the other party.[3] Another implication is that there is no negative outcome associated with lying per se. This assumption is very useful in many economic models. In contract theory it is assumed that without an explicit contract, the sides will not fulfill their respective obligations. For example, Akerlof's (1970) paper on asymmetric information and

[1] Remarks by the President Bush about Enron, at Malcolm Baldrige National Quality Award Ceremony, Washington Hilton Hotel, March 7, 2002.

[2] Important deviations from this assumption in economic modeling are Arrow's (1972) discussion of trust, Becker's (1976) modeling of altruistic preferences, and Akerlof's (1982) study of the fair-wage hypothesis. For a general discussions see Becker (1993) directions "…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). More recent models relaxing this assumption are discussed below.

[3] Note that this does not mean that a completely selfish person will always lie. There might be strategic reasons not to lie. For example, see the Kreps and Wilson (1982) discussion of reputation and imperfect information in addition to the related experiment reported in Camerer and Weigelt (1988).

the market for lemons assumed that sellers of used cars would always lie if it is to their benefit. In the literature on tax evasion, the choice of whether to avoid paying taxes is considered a decision under uncertainty; cost is treated as the product of being caught and punished, whereas benefit is simply the money saved by nonpayment. There is no cost associated with the very act of lying (Alingham and Sandmo, 1972). Another example is the game theoretic treatment of "cheap talk" (see Rabin, 1996 for a survey).

An intermediate approach is taken by utilitarian philosophers (e.g., Bentham, 1789). According to utilitarianism, in choosing whether to lie, we weigh benefits against harm and happiness against unhappiness. Similar 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 utility (not necessarily the same monetary payoffs) are equivalent (Bok, 1978, chapter 4). This is a consequentialist approach. An alternative approach that distinguishes between two decisions with the same payoff set according to the process leading to the outcomes is called non-consequentialist.

St. Augustine's approach is normative in the sense of "this is what a person should do."[4] His injunction is (unfortunately?) not supported by casual observation of real life: People do lie. Economic theory is normative in the sense of "this is 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, absent any strategic justification for doing so. The utilitarian approach predicts that if people do care about the well being of others, the decision whether to lie or not may depend on the cost to the other side.[5] As a result people will not be at the

[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 the household unity, are regarded as exceptions (Jacobs, 1960).

[5] For empirical support of the assumption that people do care about other's payoff see Guth, Schmittberger and Schwarze (1982), Kahneman, Knetsch and Thaler (1986), Dawes and Thaler (1988), Rabin (1993), and Andreoni (1990, 1995).

extreme of either always lying or always telling the truth. However, as will be shown below, people do distinguish between decisions even when the choices do not defer in 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 " innocent" choice. Hence the consequentialism assumption of utilitarian is rejected.

I start the investigation by empirically studying the role of consequences in the decision whether to lie.[6] I consider a two-person interaction in which lying increases the payoffs for the liar at the expense of her counterpart. 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 are also sensitive to how harmful lying is to the other side. Distributional models do not capture this effect although they assume inequality aversion (people are unreceptive to others earning more money than they do). This "envy" assumption is almost a polar opposite of the "aversion to doing harm" observed in the deception environment: The decision maker prefers not to lie when a lie increases her payoff a bit yet reduces the other's payoff by a great deal.[7]

To structure this observation, I propose a form of deception preferences in which a potential liar is assumed to be sensitive to the increase in her payoff associated with a lie but also to the decrease in payoff to the other side. Preferences of this type organize the data well, and for a given environment, predict the changes in the propensity of people to lie when relative and absolute payoffs are altered.

---

[6] Many other aspects of deception are studied in the literature. Psychologists study personality characteristic of honesty (Hartshorne and May, 1928), how to detect lies (Ekman, 1992, Vrij, 2001), how children learn to lie (Sodian, 1991), etc. See Ford (1995) for an introduction to the psychology of deceit, and DePaulo, Kashy, KirKendol, Wyer, and Epstein (1996) for a taxonomy of lies and their classifications according to content, motivation and magnitude. In accounting, Evans, Hannan, Krishnan, and Moser (2001) examine how preferences for wealth and honesty affect managerial reporting (see also the discussions by Baiman and Lewis, 1989, and Koford and Penno, 1992). In business, research focuses on deception in negotiation (e.g., Schweitzer and Croson, 1999, and Robinson, Lewicki, and Donahue, 2000). For what I find the most thoughtful modern treatment on the morality of deception, see Bok (1978), and particularly chapter 4.

[7] As will be discussed below assuming altruism or "matching intentions" does not work here either.

Finally, I show that cost-benefit analyses of consequences of deception are subjective, that is, they do not depend solely on actual payoffs. For example, employer-employee relations are generally marked by asymmetry in wealth. An employee may feel that her $100 gain has more "impact" than her employer's $1,000 loss. Moreover, since the judgment is subjective, self-serving biases may come into play, which make deception more likely an option.[8]

2. Definition and classification of deception and lies

It is interesting to note that the literature offers many definitions of deception and lies. These can differ, for example, by whether the liar's intentions are considered.[9] Mitchell (1986) defines deception as "a false communication that tends to benefit the communicator." This definition implies that not only humans lie but plants do too. Note that this definition assumes that the lie increases the liar's payoff. This definition is problematic because it implies that unconsciously and mistakenly misleading others is also deception. Krauss (1981) defines deception as "an act that is intended to foster in another person a belief or understanding which the deceiver consider to be false." This definition ignores a different aspect of deception, highlighted by Ekman (1992): people lie only when they do not inform others about their intention to lie in advance. Magicians, therefore, do not lie. Ekman (1992) defines deception as "a deliberate choice to mislead a target without giving any notification of the intent to do so." Vrij (2001) claims that this definition is incomplete, since liars sometimes fail to mislead. Hence Vrij defines deception as "a successful or unsuccessful deliberate attempt, without forewarning, to create in another a belief which the communicator considers to be untrue."

[8] See Babcock, Loewenstein, Issacharoff and Camerer (1995) and Babcock, Wang, and Loewenstein (1996).
[9] For a more extensive survey of modern definitions see Vrij (2001), and for classical definitions see Bok (1978).

Similarly, there are many ways to classify lies.[10] In this paper I propose to base my classification on the lie's consequences. 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 consider lies that help the other person even if it harms the liar. The motivation for this lie may be pure altruism (Becker 1976), an impure motive according to which people enjoy the act of giving (Andreoni, 1990), or an "efficiency motive," according to which people prefer outcomes that enlarge total surplus (Charness and Rabin, 2001). In the third category a lie may not help the liar but can harm both sides or at least the other part. The motive for this might be a spiteful reaction to an unfair behavior by the other side, for example.

The fourth category I consider includes lies that increase the payoff to the liar and decrease the payoff to the other side. I would argue that this is the relevant category for many economic issues, such as contract theory. If person A signs a contract with Person B, it is simply to prevent B from acting in ways that will increase her payoff at the expense of A's payoff.

In this paper I focus on the fourth category of lies, those that help the liar at the expense of the counterpart. The discussion will be devoted to the influence of different consequences of this kind of lying on the decision whether to lie. Hence, the working definition of deception I use in this paper, (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 payoff of the communicator at the expense of the other side."

[10] See Lewicki (1983), DePaulo, Kashy, KirKendol, Wyer, and Epstein (1996), and Schweitzer (2001).

3. Empirical findings

The purpose of the current paper is to study *empirically* the role of consequences on behavior in the case where a lie increases the payoff to the liar at the expense of the other side. To this aim I conducted an experiment in which I altered the absolute and relative consequences of lies, and then measured the effect of the change on the subjects' propensity to lie. Experiments are important because they allow us to observe actual behavior under conditions where this behavior has real monetary consequences for the participants. This is the reason why experiments with real monetary incentives are major tools for collecting data on economic behavior.

In addition to using experiments with real monetary incentives, it is my belief that finding out what people say about the importance of costs to the receiving end on the fairness of a lie is interesting. To examine this issue I constructed questionnaires involving "real life" scenarios. This method is well established in psychology, and has been used in economics (e.g., Kahneman, Knetsch and Thaler 1986 study of fairness).

I now turn to a description of the empirical findings.

3.1 The experiment

The Salesperson game

You wish to buy a new vacuum cleaner. There are two vacuums you are considering. Consulting with a salesperson she recommends one "because it is better." Should you follow the advice? Maybe she is paid a higher commission for this product, and that is why she recommends it? Similarly, imagine that you go to a tourist's restaurant and the waiter recommends the lobster. Are you going to follow her recommendation? If you think the waiter recommended it because she thinks it is better, you would. But what if she

recommended it simply because the cook told her that the lobster is not fresh, and must be sold today?

When are the salesperson and the waiter more likely to lie?

Procedure

The participants were 450 undergraduate students at the Technion and the University of Haifa. Students volunteered to participate in the experiment after class. They were told that the experiment will take about 15 minutes, and that everyone would be paid. In the instructions (see Appendix 1) for both player 1 and player 2, it was written that there are 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 each option. The rule were made known to both participants. After reading the instructions, player 1 was asked to send one of two possible messages to player 2:

Message 1: "Option A will earn you more money than option B."
Or
Message 2: ""Option B will earn you more money than option A."

This message was the only information player 2 had about the payoffs.

In all three treatments payoffs were constructed such that if option A was chosen, player 1 earned less than if option B was chosen, and the reverse for player 2 (recall that player 2 did not know that the payoffs were inverse). As a result, message 2 was not true, and the

payoff associated with it for player 1 was larger than the payoffs associated with message 1.[11] Actual payoffs used in the experiment are presented in Table 1 (75 pairs per cell).

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

Table 1: The different payoffs used in the experiment.

Prediction with the assumption of rational expectations and selfishness

This is a game of asymmetric information in which we have no knowledge of the beliefs of player 2; hence finding the game-theoretic equilibria with selfish players is not trivial. For the purpose of this discussion I analyze the game as a decision problem for player 1 in the following sense: I assume that player 1 has rational expectations regarding player 2's decision. That is, player 1 guesses correctly the action of player 2 given the respective message. If that is the case, and player 1 is selfish, then she will always choose the outcome that maximizes her expected payoff. As it turned out, 78% of the participants in the role of

[11] That is, "Message 1 = Truth" and "Message 2= Lie."

player 2 followed the message sent by player 1 and chose the option in which player 1 told them they would earn more money. That is, player 2 chose the option "recommended" by player 1's message. Moreover, 50 participants in the role of player 1 were asked to guess player 2's choice. Of these, 41 (82%) said that they expected player 2 to follow the message sent by them.[12]

To further test this assumption, the treatment was repeated with another group of 50 participants who played the role of player 1. After making their choices, they were told that we had already conducted the experiment with player 2 (the original instructions were adapted such that this would not contradict what they had been told previously). They were told that the player 2 they were matched with has chosen to follow the message they had sent. They were then asked whether they wished to reconsider their previous choice. Three (6%) 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 player 1 is simply interested in maximizing her own payoffs, and she has rational expectations about the reaction of player 2 to the message she sends, she should always lie. Furthermore, player 1 understands this. I find this property compelling because it helps separate strategic motives from fairness motives. Because player 1 expects the lie to "work," she is concerned only with the fairness of lying.

Results

The results of the experiment, in terms of the fraction of players 1 who lied, are presented in Figure 1. The figure is constructed such that the gain from lying for player 1 (Δx) is plotted on the horizontal axis of Figure 1, and the gain for lying for player 2 (Δy) on

[12] They were paid for accuracy according to the incentive scheme described in Dufwenberg and Gneezy (2000).

the vertical axis (note that the gain for player 2 is negative). This kind of presentation will be the basis for the preference formulation below.

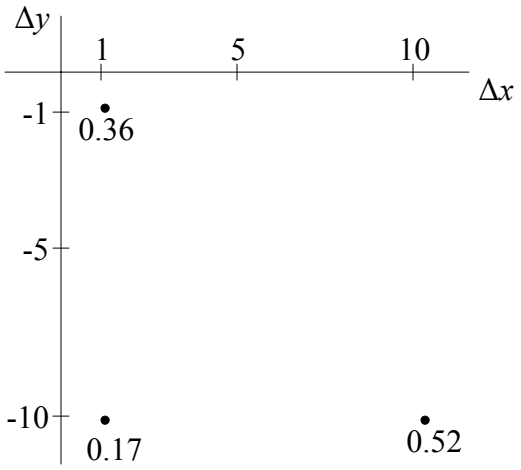

Figure 1: Fraction of participants who lied. Gains from lying for player 1 are plotted on the horizontal axis, gains from lying for player 2 on the vertical axis.

In treatment 1, where the gain for player 1 was 1 and the loss for player 2 was also 1, 27 (36%) out of the 75 participants lied. In treatment 2 (1, -10), the number who lied declined to 13 (17%) out of 75, and in treatment 3 (10,-10) it rose to 39 (52%) out of 75 participants. A statistical comparison of these differences (the *p*-values are approximated to two decimal places and calculated from the test of the equality of proportions, using normal approximation to the binomial distribution), shows that they are all significant: For the comparison of treatment 1 and 2, Z=2.64, and p=.004. For treatment 1 versus 3, Z=1.97, and p=.024 , and for treatment 2 versus treatment 3, Z=4.48 and p=.001.

Note that the results are inconsistent with those of models that assume that people are either honest or not. For example, in Koford and Penno type models (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. Their model cannot explain the outcome of the above three treatments, which indicate that people are sensitive to payoffs associated with unethical behavior and cannot be simply categorized as one of two types. Nor can those model assuming that the decision maker makes a simple cost-benefit analysis of her own monetary payoffs prior to deciding whether to lie explain the results. For example, Baiman and Lewis’ (1989) threshold model assumes that because individuals experience a small fixed disutility from lying, they are honest for all payoffs less than 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 between treatment 1 and 2, where the cost for the decision maker is fixed.

3.2 The questionnaires

The results of the experiment show that when making choices with real incentives in the laboratory, participants react differently when we change the absolute and relative consequences of their choices. But what do people think about the role of consequences in lying, and what do they say about the fairness of different lies? I study these issues with a set of questionnaires whose items referred to an empirically likely scenario. In this study, participants were students at the University of Chicago who had volunteered to fill in the questionnaires, and were paid $1 for their participation. The first group was asked the following:

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

Thus, there was no difference between the two scenarios in term of the seller's payoffs. On the other hand, the buyer's cost increased from $250 to $1,000. I used both a between-subjects design (i.e., each subject was confronted with one question, until the line saying "what would be…"), with N=50 students answering each question, and a within-subject design (i.e., the subjects answered the question for both parameters as it presented above), with N=50. The students' responses are presented in Figure 2a and 2b.

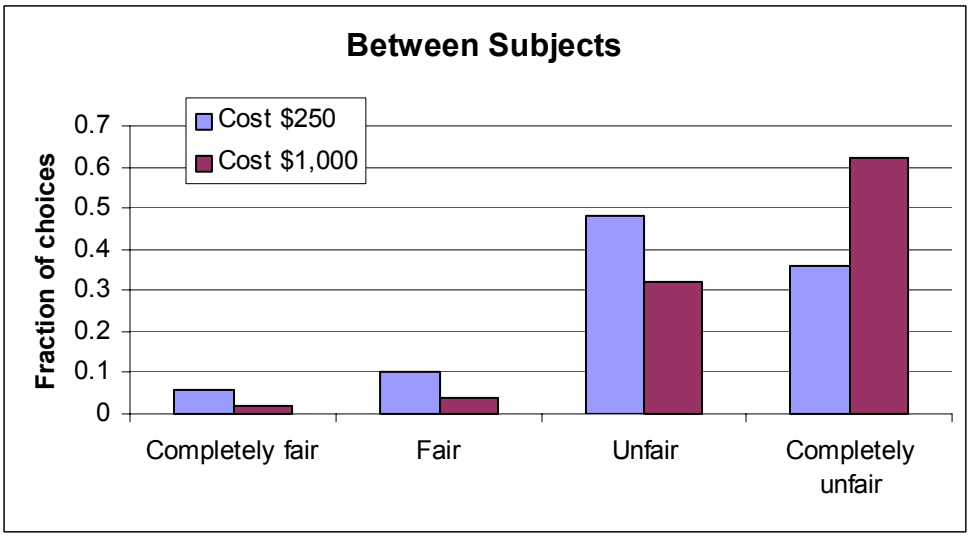

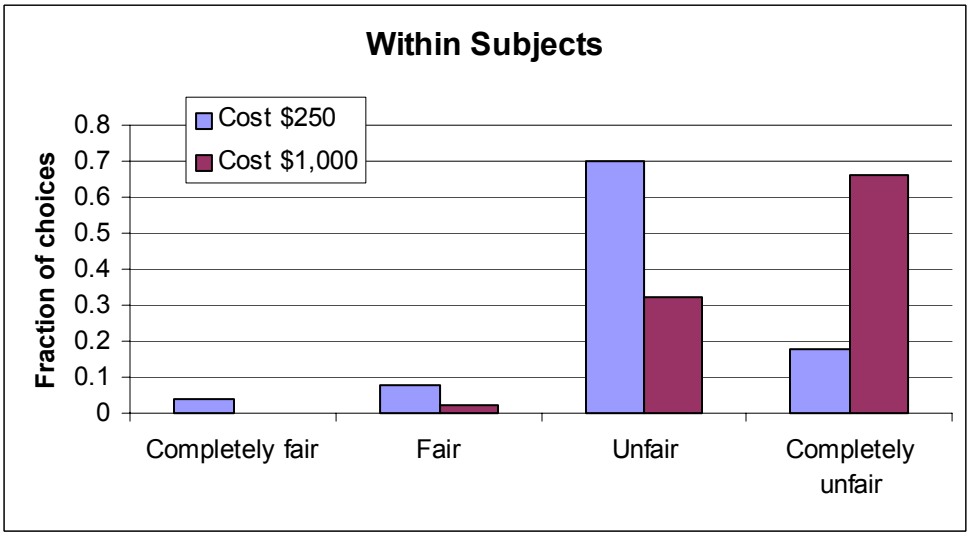

Figure 2a and 2b: Replies to the car sale question. Purchase price is $1,200, seller cost of repair is $250 and buyer cost of repair either $250 or $1,000.

The difference between the answers to the first and second question in the between-subjects design is significant ($p<.05$).[13] Inspection of the within subjects design shows a large difference in choices: In the $250 cost question, 70% of the participants chose "unfair" and 18% chose "very unfair." In the $1,000 question, only 32% chose "unfair," but 66% chose

[13] In all questions I use both Kolmogorov-Smirnov and Wilcoxson rank-sum tests.

"very unfair." This difference is highly significant (p <0.001).[14] Table 2 presents the actual choices made by each participant; as can be seen, 30 (60%) out of the 50 participants indicated that the lie was less fair when the cost was higher.

Subject #	Choice in q1	Choice in q2
1	1	2
2	1	4
3	2	3
4	2	3
5	2	3
6	2	3
7	3	3
8	3	3
9	3	3
10	3	3
11	3	3
12	3	3
13	3	3
14	3	3
15	3	3
16	3	3
17	3	3
18	3	4
19	3	4
20	3	4
21	3	4
22	3	4
23	3	4
24	3	4
25	3	4

Subject #	Choice in q1	Choice in q2
26	3	4
27	3	4
28	3	4
29	3	4
30	3	4
31	3	4
32	3	4
33	3	4
34	3	4
35	3	4
36	3	4
37	3	4
38	3	4
39	3	4
40	3	4
41	3	4
42	4	4
43	4	4
44	4	4
45	4	4
46	4	4
47	4	4
48	4	4
49	4	4
50	4	4
Average	3.02	3.66

Table 2: Reply to the car-selling question. Buying price $1,200, seller cost of fixing the pump $250 and buyer cost of either $250 or $1,000.

[14] 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 emphasis 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% chose "unfair" to the first questions, while in the between-subjects design only 48% chose this response.

The basic finding here is that people think it is less fair to lie the higher the cost for the other side. This intuition can be strengthened if we consider a third scenario in which the problem lies with the brakes, and the risk for the buyer is an automobile accident.[15]

To establish the robustness of the findings from the car sale scenario, a second set of questions regarding purchasing a house was transmitted (N=50).[16]

Mr. Johnson is about to close a deal on the sale of his house for $250,000. The roof of the house is in poor condition, and Mr. Johnson knows that if the buyer learns of this, he will have to reduce the price by $5,000 (the cost of fixing the roof). If Mr. Johnson doesn't tell the buyer, the roof will leak on the first rainy day: The resulting damage for the buyer will be $5,000. It is summer, and the only way the buyer can learn about this now is if Mr. Johnson were to tell him. Otherwise, the buyer will learn about the roof only on the next rainy day, expected only in two months.

Mr. Johnson chose not to tell the buyer about the problems with the roof.

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 $25,000 instead of $5,000?

[15] When this question was asked of 20 subjects at the University of Chicago, they found it insulting, and unworthy of an answer.

[16] Since this is a robustness test for the previous question, I used only a within-subjects design.

As can be seen in Figure 3, similar to the car sale question, increasing the buyer's cost of fixing the damage resulted in a significant decrease (p <0.001) in the perception of the lie's fairness.

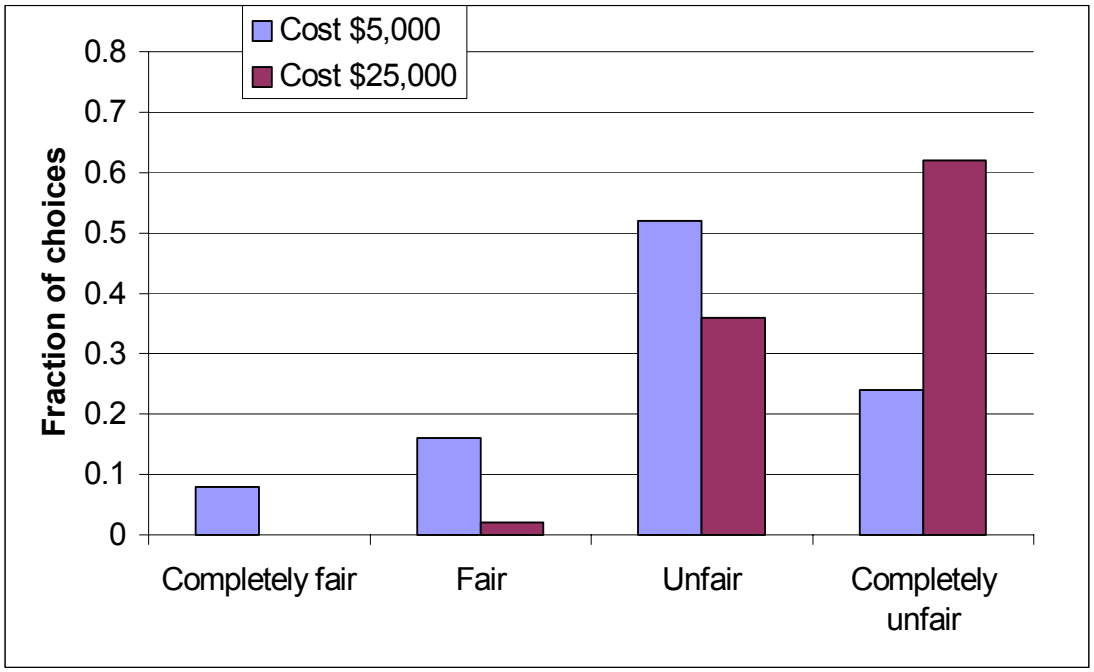

Figure 3: Reply to the house sale question. Purchase price is $250,000, seller's cost of fixing the roof is $5,000 and buyer's cost is either $5,000 or $25,000.

As can be seen from Table 3, 28 (56%) of the participants indicated that the increase in the buyer's cost makes lying more unfair.

Subject #	Choice in q1	Choice in q2
1	1	2
2	1	3
3	1	4
4	1	4
5	2	3
6	2	3
7	2	3
8	2	3
9	2	3
10	2	3
11	2	3
12	2	4
13	3	3
14	3	3
15	3	3
16	3	3
17	3	3
18	3	3
19	3	3
20	3	3
21	3	3
22	3	3
23	3	4
24	3	4
25	3	4
26	3	4
27	3	4
28	3	4
29	3	4
30	3	4
31	3	4
32	3	4
33	3	4
34	3	4
35	3	4
36	3	4
37	3	4
38	3	4
39	4	4
40	4	4
41	4	4
42	4	4
43	4	4
44	4	4
45	4	4
46	4	4
47	4	4
48	4	4
49	4	4
50	4	4
Average	2.92	3.60

Table 3: Reply to the house sale question. Purchase price $250,000, seller's cost of fixing the roof is $5,000; buyer's cost is either $5,000 or $25,000.

4. Formulation of deception preferences

Based on these empirical findings I now turn to formulate the preferences for deception.

4.1 *Related literature on fairness:*

A large body of empirical evidence has accumulated against the assumption of complete selfishness.[17] To accommodate these findings, researchers have developed formal models of social preferences that assume people are self-interested, but are also concerned about the payoffs to others. The game theoretic models of fairness can be classified into two prominent classes: models that focus on distributional concerns, and models that focus on intentions. In the distributional models, agent's preferences are influenced by the final distribution of payoffs. For example, in Kirchsteiger (1992), Fehr and Schmidt (1999) and Bolton and Ockenfels (2000), it is assumed that for a given own material payoff, a person's utility decreases with the difference between the own-payoff and that of the counterpart. Intention models (see Rabin, 1993, and Dufwenberg and Kirchsteiger, 2001) assume that intentions play a crucial role when individuals are motivated by reciprocity considerations: When considering the final distribution, the process that led to that distribution is important. Charness and Rabin (2001), and Falk and Fischbacher (1998) have developed theories that combine elements of both approaches.

The experimental game describe in this paper is similar to the dictator game in the sense that player 1 has rational expectations that player 2 will most likely follow the message she sends. Hence player 1 basically decides on the payoff distribution for the two players. The distribution models seem more adequate for this game than the intention models because the discussion is restricted to situations in which the second mover is a "dummy," leaving no place for reciprocity or intentions matching by player 2.

However, when one tries to use a distribution model in the above game, the model fails to predict the observed behavior. For example, consider a case in which player 1 has to choose between {\$6 for herself and \$5 for the other}, and {\$5 for herself and \$15 for the

[17] See for example Guth, Schmittberger and Schwarze (1982), Kahneman, Knetsch and Thaler (1986), Dawes and Thaler (1988), Rabin (1993), and Andreoni (1990, 1995). For a theoretical discussion, see Becker (1993).

other}. According to the Fehr and Schmidt (1999) model, if letting x be the decision maker's payoff and y the other person's payoff, then the decision maker's utility U is:

$$U = x - \alpha \max\{0, y - x\} - b \max\{0, x - y\}$$

Where a and b are restricted such that $0 \leq b < 1$, $b \leq a$. We get:

$U(6, 5) = 6\text{-}b$

$U(5, 15) = 5\text{-}10a$

and $U(6, 5) - U(5, 15) = 1\text{-}b\text{+}10a$

This last term is always positive given the above restrictions, meaning that the model predicts the choice of {6, 5} over {5, 15}. Moreover, keeping all else constant, the more player 2 earns in the second option, the smaller is the probability that player 1 will choose it! This seems like a good result when modeling envy, but not deception. I believe the reason for this is that a simple comparison of relative payoffs misses an important aspect of the deception problem, namely *the distinction between doing harm and inequality aversion*. In other words, people care not only about relative final outcomes, they also care about the harm done by their choices.

The intention models described above assume that people care whether the *other* player was "nice" to them or not. The argument I offer here is different since in "harm doing" environments, people also care about the role of *their own* intentions in reaching the final distribution of payoffs. To summarize, fairness models that proved highly useful in explaining/predicting behavior in distributive games based on the assumption of inequality

aversion or a preference for matching the other side's behavior, cannot be directly used to investigate deception and doing harm.

It is also the case that the standard treatment of aversion to harm doing as altruism cannot capture the phenomenon. Andreoni's (1995) compares experiments on privately provided public goods, where people tend to cooperate, with experiments on oligopolies, where a Nash equilibrium is usually observed. By constructing an experiment with different frames but the same payoff space, Andreoni (1995) is able to reject the hypothesis of pure altruism in favor of impure altruism in which "warm-glow" or the utility from the act of doing good motivates behavior. Furthermore, Andreoni (1995, p.11) states that "…there is no theory of the disutility individuals may get from the act of doing bad…" He notes the asymmetry between the marginal utility of helping when help is framed as a positive externality, and when it is framed as a negative externality.

To test the claim that harm doing behavior is different from altruism (pure or impure), I conducted a simple dictator game experiment in which a player had to choose between the monetary options described in treatment 1 in the above experiment ({6, 5} or {5, 6}). Only 34% of the participants (as compared with 64% who chose to send the true message in the deception game) chose the latter option. From these results I conclude that it is not only care for others that motivates behavior, but also aversion to lie. Note that this shows that the utilitarian assumption about consequentialism is rejected. People's choice reflects non-consequentialist preferences since they treat the choice between (5, 6) and (6, 5) differently depending on the process (a simple choice or a lie) leading to it.

4.2 *Preference for deception:*

In order to incorporate aversion to doing harm in a simple formulation, I assume that people have preferences over the consequences of their actions. Unlike the fairness models

described above, the preferences I assume relate to the consequences of choices, not final outcomes. That is, I am not looking at the final distribution of payoffs, but rather at the potential changes in payoffs resulting from a lie. This formulation involves a decision maker who has to choose whether to lie. The truth-telling payoffs are used as a reference level against which the decision maker compares the payoffs associated with the lie. This use of reference dependent utility has another nice feature: In real world trade (unlike controlled experiments), the absolute value of a deal to the other side is usually unknown, meaning that at the decision (actual signing) stage, final outcomes to the other side are unknown. For example, in the car sale scenario above, what is the absolute value of the deal to the buyer? We know how much she pays for the car, but that is not her payoff, since we do not know her utility from owning the car. The cost of fixing the car is much more likely to be known to the seller.

Consider a two-player deception game in which player 1 has to choose whether to deceive player 2. There is no strategic uncertainty. The outcome associated with the lie is an increase in payoffs to player 1 and a decrease in payoffs to player 2, relative to telling the truth.[18] The following formulation is suggested to represent player 1's preferences over the two options: Player 1 has the choice between T ("tell the truth") and L ("lie"). Let X^T be the payoffs to player 1 from choosing T, X^L from choosing L, and define Y^T and Y^L similarly for player 2. The resulting income allocations of player 1's decision between T and L are (X^T, Y^T) and (X^L, Y^L), respectively. Action L is preferred over T, if and only if:

$$U_1 = \alpha (X^L - X^T) + (1-\alpha) (Y^L - Y^T)/(X^L - X^T) > 0$$

or equivalently

[18] As discussed above, I believe that the different motivations for different lies (e.g., lies that help one at the expense of another versus lies that help both sides) result in different behavioral rules, and should not be treated in the same formulation.

$$\alpha X^L + (1-\alpha) Y^L/(X^L - X^T) > \alpha X^T + (1-\alpha) Y^T/(X^L - X^T)$$

where α $(0 \leq \alpha \leq 1)$ denotes the relative weight player 1 places on her own payoff. To simplify the presentation, define Δx as the difference in payoffs to player 1 between choosing L and choosing T, $(X^L - X^T) = \Delta x$ and to player 2 by $(Y^L - Y^T) = \Delta y$ (I assume Δx to be strictly positive and Δy negative). The utility of the decision-maker from lying is then

$$U_1 = \alpha\Delta x + (1-\alpha)(\frac{\Delta y}{\Delta x}) > 0$$

Whenever this term is 0, the decision maker is indifferent between lying and telling the truth; for positive values of U_1, the decision maker would lie. The preferences of a player with α=.5 are plotted in Figure 4. In this figure, each line represents U_1 for a given gain for player 1. As can be seen, U_1 is decreasing in the cost for player 2.

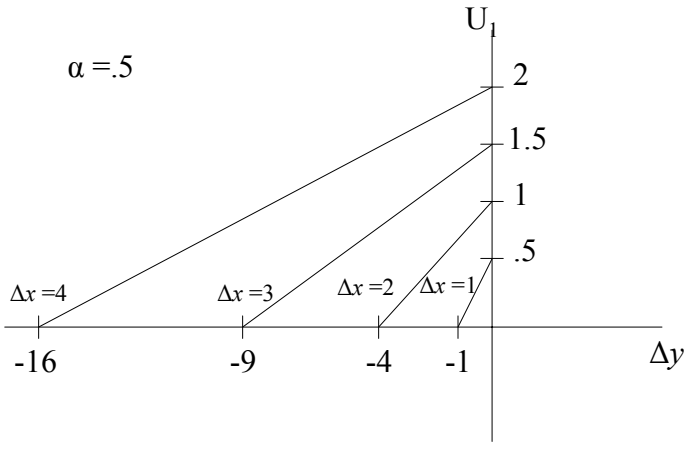

Figure 4: U_1 for a player with α=.5

Of special interest in Figure 4 is the intersection of the lines with the horizontal axis: The points at which player 1 is assumed to be indifferent between lying and not lying. That is, at these points the utility from lying equals the utility of telling the truth. For larger costs for player 2, the utility from lying is smaller than the utility from not lying, while the reverse holds for lower costs to player 2. The equation of these points is a quadratic function in the form:

$$-\frac{\alpha}{1-\alpha} \times \Delta x^2 = \Delta y$$

These points, for α=.5, are plotted in Figure 5. The equation is simply player 1's indifference curve. For values above the curve she will lie, for values below the curve she will not lie. Note that the quadratic function implies that the higher the gain for player 1, the less she cares about the relative cost for player 2.[19]

[19] This assumption is similar to Rabin's (1993, p.1291) assumption that "I now turn to the case in which material payoffs are very large. Proposition 5 states essentially that as material payoffs become large, the players' behavior is dominated by material self-interest. In particular, players will play only Nash equilibria if the scale of payoffs is large enough." See Nelson (2001) for a discussion of what happens when stakes become "very very high."

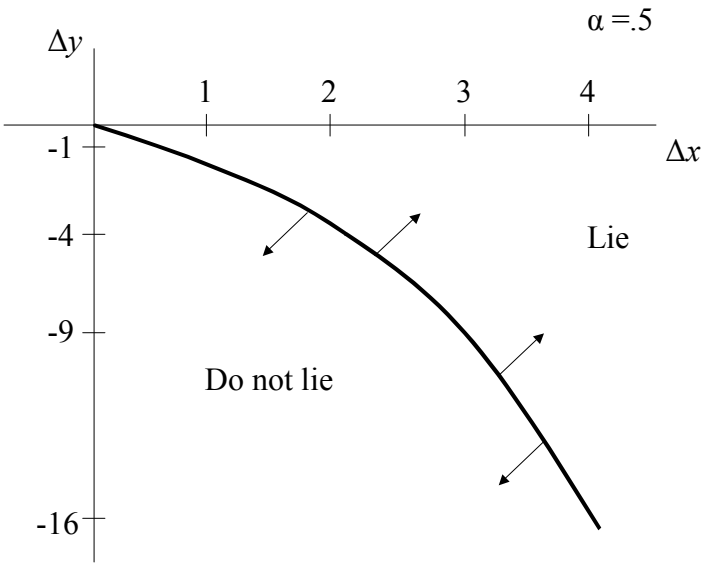

Figure 5: The indifference curve for α=.5

4.3 *Heterogeneity*

How much one cares about others differs between people. This can be described by changes in α. Extreme cases are those people who are absolutely unwilling to lie (α = o; e.g. St Augustine), and completely selfish people (α = 1 ; e.g., homo economicus). The indifference curves for α = .25, .5, and .75 are plotted in Figure 6.

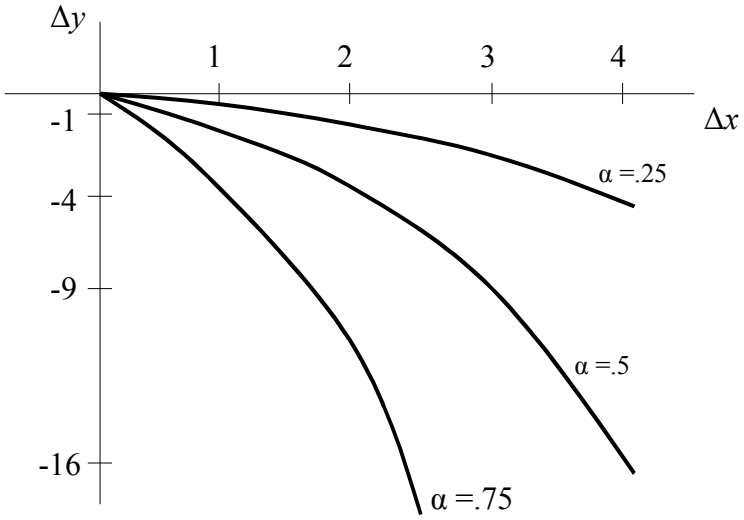

Figure 6: Indifference curves for α = .25, .5, and .75

The experiment with the game allows us to ascertain how many indifference curves were crossed when moving from one point to the other. This will be done in the next section. However, a word of caution is in place. Is it sensible to assume that each person will display the same α in all situations? In one of the classical studies in social psychology, Hartshorne and May (1928, see Ross and Nisbett, 1991, for a more general discussion), examined the honesty of school children in classroom and real-life situations. For example, they tested the willingness of children to lie to avoid getting another child into trouble, and their willingness to cheat by adding false scores to a classroom test where detection seemed impossible. Many of the specific behaviors studied were examined more than once. For example, the researchers measured children's willingness to cheat on each of several similar classroom tests. The average correlation Hartshorne and May (1928) obtained between any one type of honesty behavior and any other type was only .23. This suggests that the parameter

describing concern for others may vary between deception situations even for the same person.

4.4 *Deception versus choices between allocations*

As mentioned, comparing a choice between the same allocations in the dictator game versus the deception game results in significantly different results. I believe the force driving this difference is player 1's intentions. Player 1 is more willing to harm player 2 when her behavior is "innocent" than when she lies. This of course does not mean that the same type of formulation cannot be used to account for choices between allocations in the dictator game. It only means that participants assign different α to different intentions. The formulation can also be generalized to choices between many options, and to allow for envy to enter the decision process. Testing this conclusion and the relevant α as well as developing a more general model of choice is behind the scope of the current paper.

5. From empirical findings to the formulation

All the studies reported above supported the conclusion that people do care about the cost they impose on their counterpart when lying. For example, in the experiment, when the profit for player 1 was 1 and the cost for player 2 was 1, 36% of the participants lied. Increasing the cost for player 2 from 1 to 10, resulted in a significant drop in lying to only 17% of the participants. This trend can be interpreted as "moving down" in Figure 6, resulting in crossing the indifference curve of 19% of the subjects. For example, consider a person with $\alpha = .75$. For $\Delta x = 1$ and $\Delta y = -1$, her utility from lying is

$$U_1 = .75 \times 1 + (1 - .75) \times (\frac{-1}{1}) = .5 > 0$$

Hence, she will lie. On the other hand, for $\Delta x = 1$ and $\Delta y = -10$, her utility from lying is

$$U_1 = .75 \times 1 + (1 - .75) \times (\frac{-10}{1}) = -1.75 < 0$$

Hence, she will not lie. The results of the experiment indicate that significantly more people had an α high enough to lie when Δy = -1 but not when Δy = -10.

A similar exercise in line with the preference formulation, shows that for Δy =-10, increasing the profit for player 1 from 1 to 10 results in a significant increase in the fraction of people that were willing to lie (from .17 to .52, i.e., the indifference curve of 35% of the participants is crossed). Finally, to justify the quadratic shape of the preferences, note that multiplying both Δx and Δy by 10 results in a significant increase in the fraction of people who are willing to lie (from .36 to .52, i.e., the indifference cure of 16% of the participants is crossed). This shows that the formulation is able to capture the general structure of behavior in the empirical studies.

6. Extensions: The relative importance of consequences

6.1 A richer counterpart

A potential liar may weight changes in payoffs relative to the wealth of both sides. For example, insurance markets have asymmetry in the wealth of the two sides. The above formulation prescribes that people are more likely to lie to insurance companies than to private people, because the damage to a private person is relatively greater than the damage to a large company. Holding the costs and benefits of fraud constant, consumers' attitudes toward the acceptability of fraud affect the fraud rate. Tennyson (1997) uses consumer surveys to reveal a wide acceptance of practices that would normally be termed fraudulent, and a general perception that such practices are commonplace when the target of the fraud is a large institution. For example, in a survey reported by the Insurance Research Council (1991), consumers were asked to rate the acceptability of a wide range of insurance fraud

activities; nearly 11% of the responders found it "almost always" acceptable to allow a doctor or a lawyer to submit medical bills for treatments not received (see Tennyson, 1997 for an extensive discussion of the findings). Would the result hold if the fraud were aimed at an individual person rather than a large and rich company?

This conjecture is also of interest in labor relationships. Employers are usually wealthier than employees. Strahlberg (2001, p. 30) tells us that "interviews are a great time to perfect your lying skills." It is more or less expected from a prospective employee to "exaggerate." But what about the employer: Is it generally acceptable for the firm's representative to lie to the candidate? The following questions were used to study this issue:

Question 1:

Mr. Johnson went for a job interview in a large bank. When asked by the bank representatives about his current wage, he lied and told them that he earns $25,000 more than he actually does.

In your opinion, Mr. Johnson's behavior is (please circle one):

Completely fair Fair Unfair Very unfair

Question 2:

Mr. Johnson went for a job interview in a large bank. When he asked the bank representatives about the wage of the person he suppose to replace, they lied and told him that he earns $25,000 less than he actually does.

In your opinion, the bank's representative behavior is (please circle one):

Completely fair Fair Unfair Very unfair

I again used both a between-subjects design (N=50 in each group) and a within-subject design (N=50). The results are presented in Figure 7a and 7b.

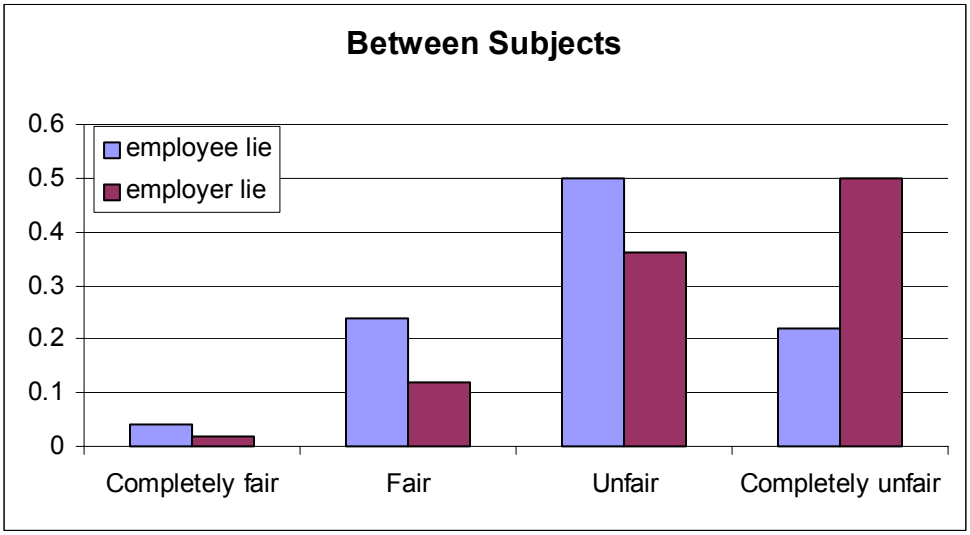

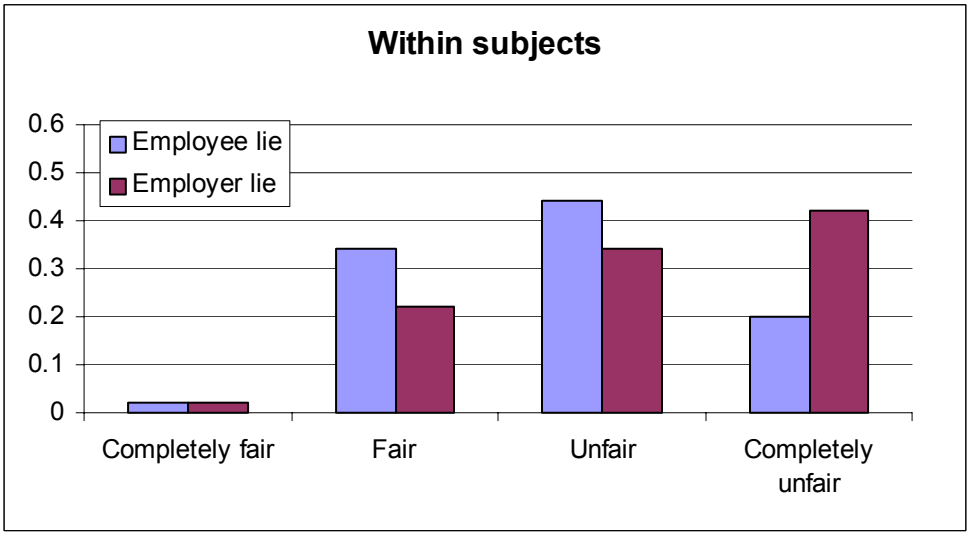

Figure 7a and 7b: Who is allowed to lie in an interview?

The difference is significant both in the between-subjects and the within-subjects investigation ($p<.01$ in both cases and both tests). Participants find lying more acceptable for the employee than for the employer (in the within-subjects treatment, 19 (38%) indicated that it is more fair when the employee lie, see Table 4). My interpretation of this result is that the

participants perceive $25,000 as much more financially meaningful to the prospective employee than to the bank; hence, it is more acceptable for her than for the bank to lie.

Subject #	Choice in q1	Choice in q2
1	1	1
2	2	2
3	2	2
4	2	2
5	2	2
6	2	2
7	2	2
8	2	2
9	2	3
10	2	3
11	2	3
12	2	3
13	2	3
14	2	3
15	2	4
16	2	4
17	2	4
18	2	4
19	3	3
20	3	3
21	3	3
22	3	3
23	3	3
24	3	3
25	3	3

Subject #	Choice in q1	Choice in q2
26	3	3
27	3	3
28	3	3
29	3	3
30	3	3
31	3	3
32	3	4
33	3	4
34	3	4
35	3	4
36	3	4
37	3	4
38	3	4
39	3	4
40	3	4
41	4	4
42	4	4
43	4	4
44	4	4
45	4	4
46	4	4
47	4	4
48	4	4
49	4	4
50	4	4
Average	2.82	3.28

Table 4: Reply to the employer-employee question.

To test for the robustness of this result, a car purchase question was used. Participants were told that the seller lied when asked about the fuel consumption of the car. One group was told that the buyer was a law student, and other participants were told that the buyer was an established lawyer. The results were replicated: Participants said that lying to the lawyer is more acceptable than lying to the law student.

A simple way to integrate this into the model is by assuming that the decision maker does not simply use Δy, but rather considers the change in utility caused by a monetary reduction of Δy. Assume for example that player 1 considers player 2's utility to be concave. In that case, the richer player 2 is (the more to the right she is on her utility function), the smaller will be the reduction in player 2's utility for a given Δy. Hence, the richer player 2 would be, the higher the chance player 1 would lie. A similar argument may be used to explain why a poorer player 1 would be more likely to lie than a richer player 1.

In many cases potential liars do not know the exact cost a lie inflicts on the other side and are required to estimate it. A series of studies (e.g., Babcock, Loewenstein, Issacharoff and Camerer, 1995, and Babcock, Wang, and Loewenstein 1996) revealed what is called a "self-serving bias" in this type of estimation. According to this bias, a psychological mechanism makes people conflate what they consider as fair with what is best for them. As a result, in an environment in which the cost of a lie to the other side is not objectively known (perhaps because the liar does not know the actual size of the cost or how wealthy the other side is), liars may undervalue that cost. If a potential liar indeed feels bad the larger the cost associated with the lie, a self-serving bias may reduce what she believes to be that cost and make her more likely to lie. This bias is called self-serving because it increases the monetary payoffs and the utility of the liar.

6.2 Extending the formulation to account for absolute payoffs

Absolute payoffs may also be important in the decision whether to lie or not. The effect of absolute payoffs on preferences may be formulated as:

$$U_1 = \alpha\Delta x + (1-\alpha)(\frac{\Delta y}{\Delta x})\frac{X^T}{Y^T} > 0$$

Adding player 1's absolute payoff from telling the truth (X^T) to the nominator makes player 1 less amenable to lying with an increase in her absolute payoff. For example, consider treatment 1 in which lying increases player 1's payoffs from \$5 to \$6. What would happen if the change were from \$105 to \$106? Adding the total payoffs to the nominator simply make the lie less likely. Adding the absolute payoff of player 2 to the denominator, results in player 1 being more likely to lie when player 2's absolute payoffs increase. For example, in treatment 2, the payoffs to player 2 are \$15 if player 1 tells the truth and \$5 if she lies. What would happen if the payoffs were \$115 and \$105, respectively? I claim that in this case, player 1 would be more amenable to lying.

Note that the comparative statics done in section 5 is still valid under this extended formulation if we consider the case in which the absolute payoffs in the case of telling the truth are fixed. Moreover, in many (most?) real-life interactions, X^T and Y^T are not observable to the decision maker (what are these values in cases of trade involving money and commodities?). In these cases, the formulation proposed in section 5 can likewise be used.

7. 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. People care how much they gain from a lie, but also about how much the other side loses (this unselfish motive diminishes with the size of the gains).

The implications of the above is 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. The formulation of preferences offered here can also explain why people are more accepting of fraudulent behavior toward large organizations or rich counterparts than towards an individual: 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.

Finally, the formulation presented in this paper can be used also for studying other aspects of harm doing, not necessarily lies. However, the findings indicate that the decision maker is not indifferent to the process leading to the outcome. That is, people have nonconsequentialism preferences in which they treat the same outcome differently depending on the process leading to it. The task of generalize the model to different scenarios of damaging others is beyond the scope of the current paper, yet I certainly view it as an interesting task.

References

Akerlof, G. A. (1970) "The market for "Lemons": Quality uncertainty and the market mechanism," *Quarterly Journal of Economics*, 488-500.

Andreoni, J. (1990) "Impure altruism and donations to public goods: A theory of warm glow giving," *Economic Journal*, 100, 464-477.

Andreoni, J. (1995) "Warm-glow versus cold-prickle: The effects of positive and negative framing on cooperation in experiments," *Quarterly Journal of Economics*, 110, 1, 1-21.

Alingham, M.G. and A. Sandmo (1972) "Income tax evasion: A theoretical analysis," *Journal of Public Economics*, 1, 323-338.

Arrow, K. J. (1972) "Gifts and exchanges," *Philosophy and Public Affairs*, I, pp. 343-362

Augustine, St. (421) *Enchiridion on Faith, Hope, and Love*. See translated by Regnery Publishing, Inc. Washington, D.C., 1996, p.21.

Babcock, L., Loewenstein, G., Issacharoff, S. and C. Camerer (1995) "Biased judgments of fairness in bargaining," *American Economic Review*, 85, 1337-1342.

Babcock, L., Wang, X., and G. Loewenstein (1996) "Choosing the wrong pond: Social comparisons that reflect a self-serving bias," *Quarterly Journal of Economics*, 111, 1-19.

Baiman, S. and B. L. Lewis (1989) "An experiment testing the behavioral equivalence of strategically equivalent employment contracts," *Journal of Accounting Research*, 27, 1, 1-20.

Becker, G.S. (1976) "Altruism, egoism, and genetic fitness: Economics and Sociobiology," *Journal of Economic Literature*, 14, 3, 817-26.

Becker, G. (1993) "Nobel lecture: The economic way of looking at behavior," *Journal of Political Economy*, 101, 3, 385-409.

Bentham, J. (1789) *The Principles of Morals and Legislation*, chap. 16, reissued 1948, Macmillan Co., Hafner Press, New York.

Bok, S. (1978) *Lying - Moral Choices in Public and Private Life*. Vintage Books, New York.

Bolton, G. and A. Ockenfels (2000), "ERC: A theory of equity, reciprocity, and competition," *American Economic Review*, 90, 166-193.

Camerer, C. and K. Weigelt (1988) "Experimental test of a sequential equilibrium reputation model," *Econometrica*, 56, 1-36.

Charness, G. and M. Rabin (2001) "Understanding social preferences with simple tests," *Quarterly Journal of Economics*, forthcoming.

Dawes, R.M., and R.H. Thaler (1988) "Anomalies: Cooperation," *Journal of Economic Perspectives*, 2, 187-197.

DePaulo, B.M., D.A. Kashy, S.E. KirKendol, A.M.Wyer, and J.A. Epstein (1996) "Lying in everyday life," *Journal of Personality and Social Psychology*, 70, 979-995.

DePaulo, B.M. and D.A. Kashy (1998) "Everyday lies in close and casual relationships," *Journal of Personality and Social Psychology*, 74, 63-79.

Dufwenberg, M. and U. Gneezy (2000) "Measuring beliefs in a lost wallet game," *Games and Economic Behavior*, 30, 163-182.

Dufwenberg, M. and G. Kirchsteiger (1998) "A theory of sequential reciprocity," mimeo.

Ekman, P. (1992) *Telling Lies: Clues to Deceit in the Marketplace, Politics, and Marriage*. Norton, New York.

Evans, J., R. Hannan, R. Krishnan, and D.V. Moser (2001) "Honesty in Managerial Reporting," *The Accounting Review*, 76, 537-559.

Falk, A. and U. Fischbacher (1998) "A theory of reciprocity," mimeo.

Fehr, E. and K. Schmidt (1999), "A theory of fairness, competition, and cooperation," *Quarterly Journal of Economics*, **114**, 817-868.

Ford, C.V. (1995) *Lies! Lies! Lies! : The Psychology of Deceit*. American Psychiatric Press, Inc. Washington, DC 20005.

Guth, W., R. Schmittberger, and B. Schwarze (1982) "An experimental analysis of ultimatum bargaining," *Journal of Economic Behavior and Organization*, 3, 367-388.

Hartshorne, H. and M.A. May (1928). *Studies in the Nature of Character, I: Studies in Deceit*. Macmillan, New York.

Insurance Research Council (1991) *Public Attitude Monitor*, Oak Brook, IL.

Jacobs, L. (1960) *Jewish Values*. Vallentine, Mitchell, London.

Kahneman, D, J.L. Knetsch, and R.H. Thaler (1986) "Fairness as a constraint on profit seeking: Entitlements in the market," *American Economic Review*, 76, 728-741.

Kant (1787) "On a supposed right to lie from benevolent motives." See *The Critique of Practical Reason and Other Writings in Moral Philosophy*, 346-50. Edited and translated by Lewis White Beck. Chicago: University of Chicago Press, 1949.

Kirchsteiger, G. (1992) *Interpersonal Comparisons and Behaviour on Markets*, PhD-dissertation, University of Vienna.

Koford, K., and M. Penno (1992) "Accounting, principal-agent theory, and self-interested behavior." In *Ethics and Agency Theory,* N. E. Bowie and R. E. Freeman (eds.), 127-142. Oxford University Press, New York.

Krauss (1981) "Impression formation, impression management, and nonverbal behaviors." In E.T. Higgins, C.P. Herman, and M.P. Zanna (eds.), *Social cognition: the Ontario Symposium*. Vol.1. 323-341, Erlbaum, Hillsdale, NJ.

Kreps, D. and R. Wilson (1982) "Reputation and imperfect information," *Journal of Economic Theory*, 27, 253-279.

Lewicki, R.J. (1983) "Lying and deception: A behavioral model," in M.H. Bazerman and R.J. Lewicki (eds.), *Negotiation in Organizations*, Sage Publication, Beverly Hills, CA.

Mitchell, R.W. (1986) "A framework for discussing deception." In R.W. Mitchell and N.S. Mogdil (eds.), *Deception: perspectives on human and nonhuman deceit*. State University of New York Press, Albany.

Nelson, W.R. (2001) "Incorporating fairness into game theory and economics: Comment," *American Economic Review*, 91, 1180-1183.

Rabin, M. (1993) "Incorporating Fairness into Game Theory and Economics," *American Economic Review*, 83, 1281-1302

Rabin, M. (1996) "Cheap Talk," *Journal of Economic Perspectives*, 10, 103-118.

Robinson, R.J., R.J. Lewicki, and E.M.Donahue (2000) "Extending and testing a five-factor model of ethical and unethical bargaining tactics: introducing the SINS scale," *Journal of Organizational Behavior*, 21, 649-664.

Ross, L, and E. Nisbett (1991). *The Person and the Situation, Perspective in Social Psychology*. McGraw-Hill, New York.

Schweitzer, M.E. (2001) "Deception in negotiation." In *Wharton on Decision Making*, S.J. Hoch, H.C. Kunreuther, and R.E. Gunther (eds.), John Wiley & Sons, New York.

Schweitzer, M.E. and R. Croson (1999) "Curtailing deception: The impact of direct questions on lies and omissions," *The International Journal of Conflict Management*, 10, 225-248.

Sodian, B. (1991) "The development of deception in young children," *British Journal of Development Psychology*, 9, 173-178.

Strahlberg, D. (2001) "Nice young man goes wild," *Chicago Business*, November 5.

Tennyson, S. (1997) "Economic institutions and individual ethics: A study of consumer attitudes toward insurance fraud," *Journal of Economic Behavior and Organization*, 32, 247-265.

Vrij, A. (2001) *Detecting Lies and Decit*, Wiley, New York.

Appendix 1: Instructions for 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.

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 the payment was 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 you will be paid 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

Instructions for 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.**

The 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 the actual payment was in the option not chosen (that is, if the message sent by your counterpart was true or not). Moreover, you will never know how much money 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