

White Lies

Sanjiv Erat and Uri Gneezy

Rady School of Management, UC San Diego

June 29, 2009

Abstract: When do people tell white lies? In this paper we distinguish between two types of white lies: those that help others at the expense of the person telling the lie that we term *Altruistic white lies*, and those that help both sides that we term *Pareto white lies*. We find that a large fraction of the participants are reluctant to tell even a Pareto white lie, demonstrating a pure lie aversion that is independent of any social preferences over outcomes. In contrast, a non-negligible fraction of participants are willing to tell an altruistic white lie that hurts them a bit but helps the other a lot. Comparing white lies to those lies where lying increases own payoff at the expense of the other reveals important insights into the interaction of incentives, lying aversion and social preferences over payoff distributions. Finally, in line with previous finding, women are less likely than men to lie when it is costly to the other side. Interestingly, we find that women are more likely to tell an altruistic lie, but tend to tell fewer Pareto ones.

1. Introduction

When people communicate they sometimes lie. Since communication is indispensable in almost any economic and social interaction, understanding when and why people choose to lie is important. A growing body of evidence suggests that the decision to lie is sensitive to incentives. People care about their own gain from lying as well as the harm that the lie may cause the other side.¹

To date, the economics literature has concentrated on what we call “Selfish Black Lies,” i.e., acts that help the liar at the expense of the other side. But sometimes lies may benefit the other person. For example, a physician may give a placebo to a patient, even if the physician believes that the substance has no specific pharmacological effect upon the condition being treated. Some moral philosophers argue against such practices even when it helps the patient at “no cost.” As St. Augustine wrote: “To me, however, it seems certain that every lie is a sin...” (Augustine, 421). Later on, philosophers like Kant (1787) have also taken this extreme approach in arguing against all types of lies, irrespective of their consequences.

Still, a recent study (Sherman and Hickner, 2008) found that 45% of doctors had given placebos in clinical practice, but only 4% of these told patients that it was a placebo. On one hand, it seems like a large fraction of doctors may be violating the moral approach cited above. On the other hand, if a doctor believes that the placebo has

¹ This literature builds on the seminal work of Crawford and Sobel (1982). For a treatment of the role of incentives in the decision to lie, see Gneezy (2005). See also Boles et al (2000); Crawford (2003); Brandts and Charness (2003); Croson et al (2003); Ellingsen and Johannesson (2004); Ellingsen, Johannesson and Lilja (2009); Charness and Dufwenberg (2006); Sánchez-Pagés And Vorsatz (2007); Mazar, Amir and Ariely (2008); Dreber and Johannesson (2008); Sutter (2009); Hurkens and Kartik (2009); Lundquist, Ellingsen, Gribbe and Johannesson (2009); Ellingsen, Johannesson, Lilja and Zetterqvist (2009).

therapeutic value (as 96% of the physicians in the study did), then why not use a placebo since it helps the patient without causing harm to anyone?² The utilitarian approach (e.g., Jeremy Bentham, 1789) argues that one *should* lie in such situations. When considering whether to lie, a utilitarian would argue, one should weigh benefits against harm, and happiness against unhappiness. The act of lying in itself carries no bad consequences. To quote Martin Luther “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.”

This kind of “useful lies” is the focus of the current paper. We wish to expand the discussion in the economics literature to “White Lies”-- lies that help the other person. In some cases, a lie can harm the liar but help the other person; we call this “Altruistic White Lies.” People may choose to tell such lies because they care about the other person’s payoffs. For example, such lies may create efficiency of the type Charness and Rabin (2003) discuss if the lie results in a smaller loss to the liar relative to the gain obtained by the other person. While such a lie is not Pareto improving, if someone cares enough about the other person and about increasing the total pie, she may choose to lie.

The second type of white lies are those that constitute a Pareto improvement, i.e., when both sides earn more as a result of the lie. We call such lies “Pareto White Lies.” Absent a cost of lying, one would expect people to always tell such lies.

Figure 1 presents our taxonomy of lies based on consequences. The dimensions we use in the figure are the *change* in payoffs resulting from a lie.

² We abstract in the example from strategic issues such as reputation that are associated with the use of placebo, e.g., “In the clinical setting, the use of a placebo without the patient’s knowledge may undermine trust, compromise the patient-physician relationship, and result in medical harm to the patient.” http://www.ama-assn.org/ama1/pub/upload/mm/369/ceja_recs_2i06.pdf.

Understanding when people choose to tell white lies is crucial to our understanding of deception. First, people who are reluctant to tell Pareto White Lies demonstrate lie aversion that is independent of social preferences over outcomes. It shows that some people refrain from lying not because of the consequences, but rather because they simply view lying as a bad act in itself. As such, this provides the best test of a pure cost of lying, in line with the moral stand.

Second, contrasting white lies with black lies helps in understanding the interaction between distributional concerns and lying aversion. It also helps in identifying different types of people. For example, one may expect that the motivation to tell a Selfish Black Lie arises from putting more weight on own payoffs. On the other hand, the motivation to tell an Altruistic White Lie may rise from having a higher weight on the other's payoffs in one's utility function.

Third, as Dreber and Johannesson (2008) nicely demonstrated, there are gender differences in the tendency to lie. In particular, they find that men are more likely to tell a Selfish Black Lie. Contrasting men's and women's behavior in different domains of lying allows us to test the gender differences in the interaction of distributional concerns and lying aversion. We first successfully replicate the results of Dreber and Johannesson (2008): Men are significantly more likely in our experiment to tell a Selfish Black Lie. Moreover, men are also more likely to tell a Pareto White Lie.

Interestingly, it is the women who are more likely than men to tell an Altruistic White Lie. This result shows an interesting interaction between lie aversion and social preferences. It appears as if women have a higher cost of lying, but at the same time are

also more sensitive to the other person's payoffs (see Croson and Gneezy, 2009 and Eckel and Grossman, 2008 for surveys of gender differences in social preferences).

2. Experimental Design and Procedure

Sender-receiver game:

There are two players acting sequentially in the role of a sender and receiver respectively. We rolled a 6-sided dice before the start of the game, and communicated the outcome only to the sender. The sender was then asked to send a message to the receiver from a pool of six possible messages. The six possible messages are "The outcome of the roll of dice was i " where $i \in \{1,2,3,4,5,6\}$. Full set of instructions are given in Appendix 1. The sender is informed that the payment in the experiment will depend on a choice made by the receiver. She is also told that the only information the receiver will have regarding the actual outcome of the roll of the dice is her message. There are two payment options, A and B. The sender knows the payoffs (to both players) associated with each option, and she is also told that the receiver will not know the payoffs. Finally, the sender is told that if the receiver chooses the real outcome of the dice roll, payment option A will be implemented. Otherwise, both will be paid according to option B (the payoffs associated with each option are described below).

The message from the sender is the only information regarding the roll of the dice that the receiver has. After observing this message, the receiver is asked to choose a number from the set $\{1,2,3,4,5,6\}$. This choice determines which of two possible payoff options, A or B, get implemented (the receiver is not told what these payoffs are). She is told, just as the sender was, that if the receiver chose the actual outcome of the roll of

dice, payoff option A is implemented, and for any other choice, payoff option B is implemented.

Previous research that has examined deception using sender-receiver games have typically utilized a smaller message set. For instance, Gneezy (2006) and Dreber and Johannesson (2008) utilized a message space with one truthful message and one deceptive message. They interpreted the fraction of deceptive messages as the fraction of senders who (attempt to) deceive the receivers. Sutter (2009) nicely demonstrates that a sender who expects the receiver to disbelieve any received message may undertake (implicit) deception by sending a truthful message. His experimental data reveals that, while the original conclusions from the aforementioned research survive this correction, some senders do engage in such sophisticated deception. To avoid this potential confound in our experiment, we use a richer message space that minimizes the likelihood of senders engaging in this type of sophisticated deception.³

Procedure:

Table 1 summarizes the payoffs in the 9 different treatments used in this article. As Table 1 illustrate, the treatments differ in the relative gains/losses to the sender or receiver when the receiver chooses a number other than the actual outcome of the roll of

³ To see why the richer message space can address the issue of sophisticated deception, consider a sender who prefers option B over option A (i.e., $u_B > u_A$). Suppose that the sender believes that with probability p the receiver will disbelieve her message and try to invert her message by choosing a number randomly from the remaining 5 messages. Then, her (expected) value from sending a truthful message is $V_T = (1 - p)u_A + pu_B$, and her value from sending a deceptive message (assuming no lying costs) is $V_D = (1 - p)u_B + p(4/5 u_B + 1/5 u_A)$. Thus, the sender will engage in sophisticated deception (i.e., $V_T > V_D$) only if $p > 80\%$. While it is still feasible in our game to have senders who believe that more than 80% of the receivers will disbelieve and invert the message, the expanded message space makes this less likely that such sophisticated deception will play a major role. In our experiment, 41 (82%) of the receivers chose a number equal to the message they received.

the dice. Note that in all but the last treatment, if the receiver chooses the actual outcome of the roll of the dice, each player gets \$20.

Treatment	N	A	B
T[-5,10]	120	(20, 20)	(15, 30)
T[-1,10]	121	(20, 20)	(19, 30)
T[1,10]	118	(20, 20)	(21, 30)
T[10,10]	114	(20, 20)	(30, 30)
T[1,0]	123	(20, 20)	(21, 20)
T[1,-1]	92	(20, 20)	(21, 19)
T[1,-5]	85	(20, 20)	(21, 15)
T[10,0]	90	(20, 20)	(30, 20)
Reverse	85	(30, 30)	(20, 20)

Table 1: Payoffs (in \$) in the different treatments. N is the number of senders per treatment.

Finally, we also compare the results of T[10,10] in which lying yields an extra \$10 to each player with a simple reversal of payoffs (“Reverse” treatment). In this treatment if the receiver chooses the correct number then the payoff is 30:30. If she chooses a wrong number the payoff is 20:20. This extra treatment serves two goals. First, as in Gneezy (2005), this comparison helps rule out a rather unlikely alternate hypothesis, that the people who do not lie in T[10,10] are truthful not because of any lie aversion, but because they actually prefer the payoff distribution of 20:20 over 30:30. In the Reverse treatment, the higher payoff distribution is also associated with telling the truth.

More importantly, the Reverse treatment can help us figure out how many of the senders expect the receiver to follow their message. If senders choose the truthful message (the one corresponding to the actual outcome) because they expect the receiver not to follow the message (as in Sutter’s 2009), then we expect them not to choose the truthful message in the Reverse treatment. In that sense, the Reverse treatment also offers

an indirect way to measure the beliefs of senders regarding the receiver's choice to follow the message or not.

The experiment was run online through two different online sites. There was no significant difference between the sites, and so we pool the data together. The participants in addition to completing the experiment also reported their gender. After cleaning the data to remove incomplete experimental sessions and duplicate IP addresses (to avoid having a single subject participate multiple times), we were left with 948 senders.

As described in the experimental instructions, 1 out of 20 senders were paired up with a receiver, and the payment to the sender-receiver pairs was awarded in Amazon gift certificates.

3. Results

The fraction of people who chose to lie, per treatment, is presented in Figure 2. The origin in the figure represents the payoffs of \$20 to each player and is obtained when the receiver chooses the actual outcome. Hence, the different treatments represent the deviations in payoffs resulting from the receiver choosing a number different than the actual outcome of the dice roll. The fractions represent the actual fraction of senders who chose to lie.

<Insert Figure 2>

We now discuss the main empirical results. The first result regards Altruistic White Lies, in which the sender loses money and receives gains money if option B is implemented.

Result 1: Some senders lie when it costs them a little but helps the receiver a lot.

A substantial fraction of the senders in the Altruistic White Lies treatments ($T[-5,10]$ and $T[-1,10]$) choose to lie. This finding lends strong support for social preferences. 29% of the senders were willing to lose \$5 in order for the receiver to gain \$10. While the payoff resulting from lying is not a Pareto improvement over telling the truth, lying does increase the total surplus. It is interesting to note that there is no statistically significant difference in the fraction of participants who lie between $T[-5,10]$ (29%) and $T[-1,10]$ (33%), despite the fact that the sender's cost goes down from \$5 in the former to only \$1 in the latter.⁴ It appears as though the people who are willing to tell Altruistic White Lies are not very sensitive to their monetary cost in doing so (at least as long as the lie results in efficiency gains).

The second result is on Pareto White Lies and lie aversion.

Result 2: Some senders do not lie even when lying results in a Pareto improvement

Even in $T[10,10]$, when each participant earns extra \$10 from a lie, only 61% of the senders lie. This finding offers strong support to the hypothesis that some people have

⁴ 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.

high costs associated with lying. This result represents clean evidence of lie aversion without distributional preferences. It would be interesting, in future research, to see how other factors influence this cost of lying. For example, in the domain of Selfish Black Lies, Lundquist, Ellingsen, Gribbe and Johannesson (2009) study lying aversion along two dimensions not studied by Gneezy (2005). First, they show that the aversion to lying is stronger the further you deviate from the truth. Second, using a richer set of messages, they allow for promises (as in Ellingsen and Johannesson, 2004 and Charness and Dufwenberg, 2006) and find that the aversion to lying depends on the strength of the promise.

As Kartik (2009) shows, lying costs or aversion are important in modeling deception. Kartik (2009) uses the classic strategic communication setting of Crawford and Sobel (1982) with the addition assumption that some messages entail exogenous or direct lying costs for the Sender (see also Matsushima, 2008 and Ottaviani and Squintani, 2006).

The results of the “Reverse” treatment lend further support to the presence of lie aversion: 94% of the participants chose to tell the truth and promote the 30:30 allocation over the 20:20 one. The difference between the fraction of people who chose to lie in T[10,10] (60%) and those who lie in the Reverse treatment (6%) is highly significant ($Z=7.99$, $p<0.001$).

Besides offering support for the lie aversion argument, the comparison of T[10,10] and the Reverse treatment shows that the high fraction of truth telling in the former is not due to what Sutter (2009) called sophisticated lying. That is, the senders apparently expected the receivers to follow the message they are sending, and this is why they chose to send a truthful message in the Reverse treatment.

Not surprisingly, more people were willing to lie in T[10,10] (61%) than in T[1,10] (46%) ($Z=2.38$, $p=0.008$). That is, some senders had a cost of lying that was high enough that they avoided sending a deceptive message when their benefit was only \$1, but not when their benefit was \$10. This finding offers support to the hypothesis that the decision to lie depends on the incentives involved. More senders are likely to lie when incentives to do so are higher.

It is also interesting to note that fewer senders lie when only one party gains from the lie relative to when both benefit from it. As identified above, significantly more people are willing to lie in T[10,10] (61%) compared to T[1,10] (46%); however, the difference between the fraction of participants who lie in T[10,0] (47%) and T[1,0] (41%), while in the right direction, is statistically insignificant. Similarly, the difference between the fraction of participants who lie in T[10,10] (61%) and T[10,0] (47%) is large and significant ($Z=2.1$, $p=0.018$); however, the difference between T[1,10] (46%) and T[1,0] (41%), while in the right direction, is not significant. That is, a sender is more sensitive to own benefits and receiver's benefits and is more likely to lie when both benefit from the lie. It seems as if senders who help the receiver while helping themselves feel better about lying than those who do not help the receiver in the process. Increasing the receiver's payoff helps justify the lie.

The third result we discuss is on Selfish Black Lies.

Result 3: Senders are sensitive to the receiver's cost associated with a lie

The difference between fraction of lies in T[1,-1] (31%) and T[1,-5] (20%) is significant ($Z=1.75$, $p=0.04$). This result is in line with other findings in the literature that for a given own benefit, senders are less likely to lie when the cost to the receiver increases (e.g., Gneezy, 2005; Dreber and Johannesson, 2008; Sutter, 2009).

More generally, Figure 2 shows an important trend. More senders lie as we move right, i.e., increase the sender's payoff. At the same time, moving up (i.e., decreasing the receiver's cost or increasing the receiver's benefit) results in an increase in the fraction of senders that lie. Our results generalize the findings regarding the importance of incentives in the decision to lie to the case of white lies, and show that senders are sensitive to their payoff (benefits or costs) as well as the receiver's payoff (benefits or costs).

This aggregate result is likely to hide individual differences. For example, it seems plausible that an "altruistic" sender may be willing to tell an Altruistic White Lie, but not to tell a Selfish Black Lie. A different "selfish" sender may have just the opposite preferences.

An important aspect of individual differences is the case of gender differences. Do women and men have different propensity to lie, and if so, how does this depend on incentives? The next section addresses this question.

Gender differences in the propensity to lie:

The literature documents systematic differences between women and men in decision making (for surveys, see Eckel and Grossman, 2008; Croson and Gneezy, 2009). Is there such a difference between men and women in their propensity to lie?

In an interesting recent study, Dreber and Johannesson (2008) use Gneezy's (2005) game and find that, in the domain of Selfish Black Lies, men are significantly more likely to lie to secure a monetary benefit. Our findings successfully replicate this result. More men tell a selfish black lie compared to women ($Z=1.33$, $p=0.091$ for $T[1,-5]$ and $Z=1.06$, $p=0.144$ for $T[1,-1]$).⁵ Can we conclude from this that men are always more likely to lie? The following result suggests that we cannot.

Result 4: Women are more likely to tell an Altruistic White Lie than men

The fraction of men and women senders who chose to lie, per treatment, is presented in Figure 3. In both $T[-5,10]$ and $T[-1,10]$, more women lied than men ($Z=1.55$, $p=0.061$ for $T[-1,10]$, and $Z=1.02$, $p=0.154$ for $T[-5,10]$). That is, unlike the case of Selfish Black Lies, women senders were significantly more willing than men to lie and endure the relatively small cost to themselves, as long as it helped the receiver more⁶.

The final result is about gender differences in the domain of Pareto White Lies.

Result 5: Women senders are less likely to tell a Pareto White Lie than men

More women seem to have a cost of lying that is higher than the benefit of the Pareto White Lies⁷ ($Z=3.391$, $p<0.001$ for $T[1,10]$, and $Z=2.009$, $p=0.022$ for $T[10,10]$).

⁵ The results of a logistic regression of propensity to lie in treatments $T[1,-5]$ and $T[1,-1]$ shown in Table 3 reveals a significant main effect of gender (Walds Chi-Square = 2.757, $p=0.097$).

⁶ The results of a logistic regression of propensity to lie in treatments $T[-5,10]$ and $T[-1,10]$ shown in Table 4 reveals a significant main effect of gender (Walds Chi-square = 3.303, $p=0.069$).

⁷ The results of a logistic regression of propensity to lie in treatments $T[1,10]$ and $T[10,10]$ shown in Table 5 reveals a significant main effect of gender (Walds Chi-square = 14.378, $p<0.000$).

For example, 56% of the women compared to 76% of the men were willing to lie when the lie increased both participants' payoffs by \$10.

Also, comparing the fraction of men and women who lie in treatments $T[-1,10]$ and $T[1,10]$ shows an interesting effect of how gender influences the sensitivity to incentives: When own payoffs change from a \$1 cost to a \$1 benefit, the fraction of men who lie increases significantly (increases from 24% to 65%, $Z=3.91$, $p<0.001$), however, the fraction of women who lie does not change significantly (changes from 38% to 33%).

4. Conclusion

Why should economists care about small lies with seemingly little economic effect? In contrast with psychologists, an important challenge to behavioral economists is to show that the finding has some impact on important economic behavior.

One example that illustrates the importance of “small” dishonest behavior is downloads versus purchases in the music industry.⁸ Evidence from this industry suggests that millions of people are willing to lie for under \$20, and in many cases for under \$1. One person who downloads music without paying engages in dishonest behavior with very little effect on the industry. Millions of people who do this change the industry profoundly. In fact, illegal downloads changed the music industry dramatically.

As in our results, it is interesting to note the co-existence of dishonest behavior and honest behavior in the same market. While many people do not pay for the music they download, many others use iTunes or buy CDs, paying hundreds of millions of dollars every year. Even people who can download for free choose to pay, despite of only a negligible risk of a penalty.

⁸ We are grateful to Eddie Dekel for suggesting this example.

This example illustrates the economic importance of understanding even small dishonest behavior. So, why do some people avoid lying? One explanation put forward by Charness and Dufwenberg (2006) (see also Dufwenberg and Gneezy, 2000; Dufwenberg and Battigalli, 2009) is that people experience belief-based guilt. Specifically, people experience a disutility when they let down others. Further, their model suggests that the size of this disutility is related to the difference between the consequence of the action and the other person's expectation of the consequence of own action. The aversion to lying in our experiment (specifically in the domain of Pareto White Lies) cannot be explained by this theory of guilt based on monetary consequences since the consequence of lying in each of these cases is a Pareto improvement. In particular, the person who is lied to also earns more money as a result of the lie. This seems to suggest that (at least a part of) the reason people do not lie may have to do with their endogenous lying cost and not guilt over the consequences. Indeed, our results do not preclude the possibility that people feel guilty over the act of lying itself. However, our results do demonstrate that the aversion to lying cannot be solely explained by the negative consequences of the lie.

A different take is proposed by Ellingsen and Johannesson (2004) who model a commitment-based guilt where the act of breaking a promise incurs an endogenous cost. In a recent study, Vanberg (2008) contrasted these two theories of guilt (expectation-based and commitment-based) and offered evidence suggesting that people's second order beliefs (i.e., beliefs about other's expectations) do not significantly affect whether someone breaks a promise or not. From this, he concludes that "people have a preference for promise keeping, per se." Still, since the study considers settings where breaking a

promise results in benefit to oneself and cost to the other, it leaves open the question as to whether the consequences affect this preference, specifically do people have a preference for keeping their promise because of the promise, or does the preference for truthfulness arise because of the negative consequences of breaking the promise. In line with Ellingsen and Johannesson (2004) model, our study suggests that such preferences, while influenced by the consequences, are present even when the consequences are positive (Pareto improving).

We do not claim that expectations are irrelevant to whether one lies or not. A person who does not “bluff” in poker is a bad player. The expectation in such a context would be that everyone lies. A related alternative is that people feel guilty when they lie and violate a social norm. The amount of guilt they feel depends on the descriptive norm, i.e., their beliefs about the adherence to the norm in their peer group. Specifically, people would feel greater guilt about lying when they expect their peer group to be more honest (see Lundquist et al, 2009 for experimental evidence).⁹

For demonstration, consider the simple model with two outcomes A and B. Suppose that a person’s utility from outcomes A and B are U_A and 0 respectively. The outcome B is obtained when the person is honest and the outcome A is obtained when the person lies (and violates the norm). A person experiences disutility c when she believes that no one in her peer group lies and she lies. When she believes that a fraction p of her peers lie, she experiences $(1-p)c$ disutility from lying. Hence, her total utility from lying is $U_A - (1-p)c$ and from being honest is 0.

⁹ An alternate explanation can focus on a person’s perception (belief) about the strength of the injunctive “do-not-lie” social norm, i.e., a person’s beliefs about what fraction of people find lying acceptable. Still, these two beliefs (the injunctive norms and the descriptive norms) are likely to be highly correlated (unless the person believes that his/her peer group are hypocrites). So for the purposes of our explanation, it seems adequate to focus only on one.

In this explanation, as in Ellingsen and Johannesson (2004), and as was demonstrated in Vanberg's (2008) experiment, the second order beliefs do not impact lying aversion. It is also consistent with Charness and Dufwenberg's (2008) results if when a person believes that her peer group is more likely to lie, she is also likely to expect others to believe so as well (false consensus effect). Thus, her second order expectation about other's beliefs about lying (as in Charness & Dufwenberg, 2008) is likely to be correlated with p , her first order beliefs about prevalence of lying in her peer group. This correlation could explain why people who had greater second order beliefs in their experiment were also more likely to lie.

Our simple explanation also predicts that one's belief about the descriptive norm in a given context affects lying aversion and offers a rationale for why people may not feel much guilt in lying in poker as people expect no one to tell the truth in such contexts. That is a person's *first order* beliefs about the descriptive norm that everyone engages in deception. An important avenue for future research includes investigation into the source of beliefs about the descriptive norm of own peer group.

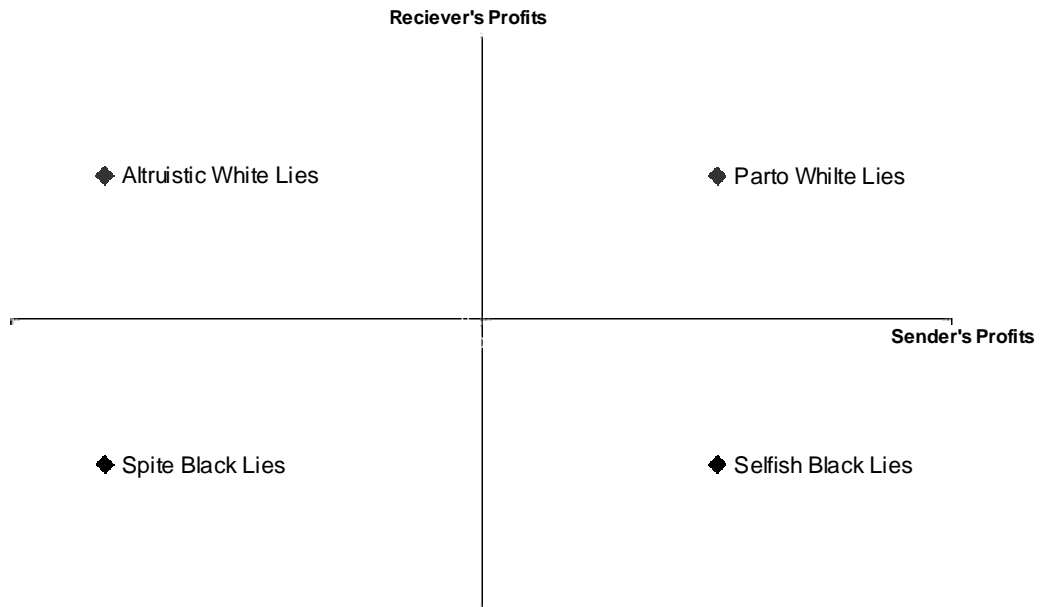


Figure 1: Taxonomy of Lies Based on Change in Payoffs

The origin in the figure represents the payoffs resulting from telling the truth. We define “White Lies” as lies that increase the other side’s payoffs, i.e., above the zero line in the horizontal dimension in the figure. “Black Lies” are those that decrease the other side’s payoffs, i.e., below the horizontal zero axes. Costly lies are defined as below zero on the vertical dimension.




Figure 2: The proportion of senders who lied in each of the eight treatments.

	T[-5,10]	T[-1,10]	T[1,10]	T[10,10]	T[1,0]	T[1,-1]	T[1,-5]	T[10,0]
T[-5,10]		-0.04 Z=-0.652 P=0.257						
T[-1,10]	0.04 Z=0.652 P=0.257		-0.13 Z=-2.01 P=0.022					
T[1,10]		0.13 Z=-2.01 P=0.022		-0.15 Z=-2.39 P=0.008	0.05 Z=0.673 P=0.251			
T[10,10]			0.15 Z=2.39 P=0.008					0.14 Z=2.1 P=0.018
T[1,0]			-0.05 Z=-0.67 P=0.251			0.09 Z=1.492 P=0.068		
T[1,-1]					-0.09 Z=-1.49 P=0.068		0.12 Z=1.746 P=0.04	
T[1,-5]						-0.12 Z=-1.75 P=0.04		
T[10,0]				-0.14 Z=-2.1 P=0.018				

Table 2: The difference in proportion of lies.

Each cell represents the difference in proportion of lies. For instance, for the cell [T[-5,10],T[-1,10]], the number -0.04 is proportion of lies in treatment T[-5,10] – proportion of lies in treatment T[-1,10]. All numbers given in bold are significant at 10% or less.

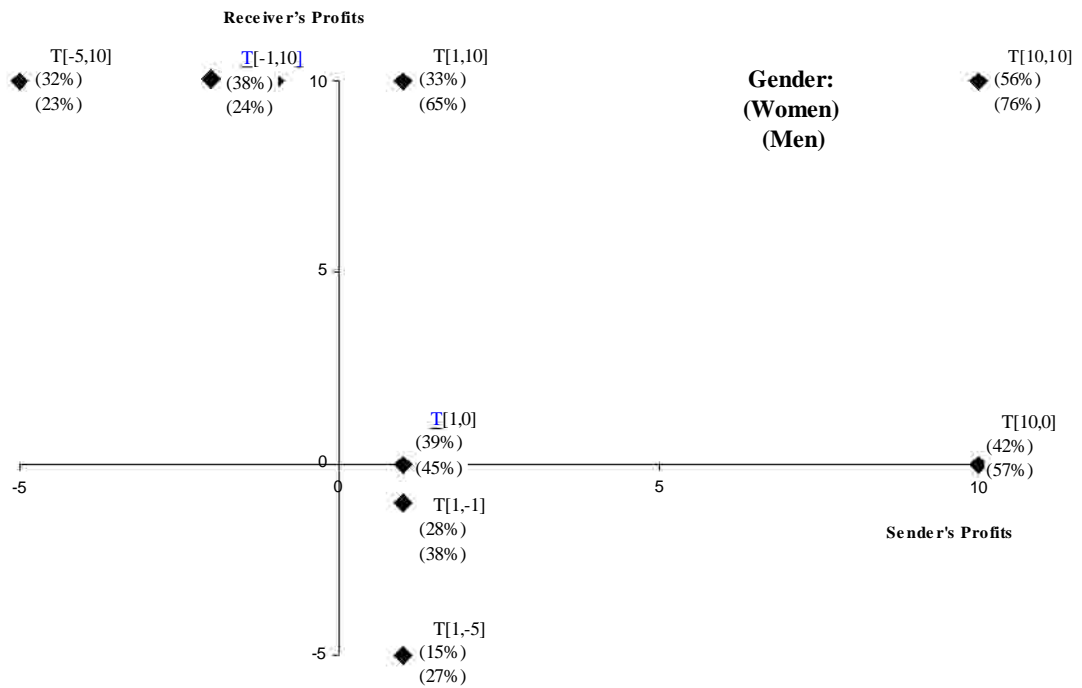


Figure 3: The proportion of women and men senders who lied in each of the eight treatments. The number on the top in each treatment is the fraction of women who lied, and number on the bottom is the fraction of men who lied.

Variable	DF	Parameter Estimate	Error	Wald Chi-Square	Pr > Chi-Square
Intercept	1	-0.422	0.305	1.907	0.167
Treatment (T[1,-5]=1)	1	-0.632	0.355	3.161	0.075
Gender (Female=1)	1	-0.584	0.352	2.757	0.097

Table 3: Logistic Regression for Selfish Black Lies (Treatments T[1,-1] and T[1,-5])

Variable	DF	Parameter Estimate	Error	Wald Chi-Square	Pr > Chi-Square
Intercept	1	-1.278	0.301	17.995	0.000
Treatment (T[-1,10]=1)	1	0.210	0.281	0.556	0.456
Gender (Female=1)	1	0.558	0.307	3.303	0.069

Table 4: Logistic Regression for Altruistic White Lies (Treatments T[-5,10] and T[-1,10])

Variable	DF	Parameter Estimate	Error	Wald Chi-Square	Pr > Chi-Square
Intercept	1	0.528	0.266	3.938	0.047
Treatment (T[10,10]=1)	1	0.797	0.281	8.028	0.005
Gender (Female=1)	1	-1.155	0.305	14.378	0.000

Table 5: Logistic Regression for Pareto White Lies (Treatments T[1,10] and T[10,10])

Appendix A

INSTRUCTIONS FOR PLAYER 1 (SENDER)

Welcome to our short online experiment. Please read these instructions carefully. You may earn a considerable sum of in Amazon Gift certificates, depending on the decisions you make in the experiment. We will randomly choose 1 participant out of 20 and pay this person, privately and in gift certificate, next week. The rest of the instructions describe what the procedure will be should you be chosen.

You will be matched randomly by the computer with another participant in this online experiment. Neither of you will know the identity of the other.

Before starting this online experiment, we have rolled a 6-sided dice, and obtained the outcome 2. The other participant does not know that the outcome was 2.

We now ask you to send a message to the other participant. The message will correspond to a number from 1 to 6. The 6 possible messages are:

Message 1: “The outcome from the roll of the dice is 1”

Message 2: “The outcome from the roll of the dice is 2”

Message 3: “The outcome from the roll of the dice is 3”

Message 4: “The outcome from the roll of the dice is 4”

Message 5: “The outcome from the roll of the dice is 5”

Message 6: “The outcome from the roll of the dice is 6”

Your message will be shown to the other participant, and then he or she will be asked to choose a number between 1 and 6. The message you send is the only information the other participant will have regarding the roll of the dice.

The choice of the number by the other participant (who does not know the outcome of the roll but will see your message before choosing) will determine the payments in the experiment.

- If he/she chooses the actual outcome of the roll of the dice, 2, then you will be paid \$20 gift certificate and he/she will be paid \$20 gift certificate.
- If he/she chooses a number different than the actual outcome, you will be paid \$X gift certificate and he/she will be paid \$Y gift certificate

The message I will send is: “The outcome from the roll of the 6-sided dice is _____”

INSTRUCTIONS FOR PLAYER 2 (RECEIVER)

Welcome to our short online experiment. Please read these instructions carefully. You may earn a considerable sum in Amazon Gift Certificates, depending on the decisions you make in the experiment.

You will be matched randomly by the computer with another participant in this online experiment. Neither of you will know the identity of the other.

Before starting this online experiment, we have rolled a 6-sided dice, and told the outcome of it to the other participant, but we are not going to tell it to you.

After being informed of the roll of the dice, the other participant has sent a message to you. The message corresponds to a number from 1 to 6. The 6 possible messages are:

Message 1: “The outcome from the roll of the dice is 1”

Message 2: “The outcome from the roll of the dice is 2”

Message 3: “The outcome from the roll of the dice is 3”

Message 4: “The outcome from the roll of the dice is 4”

Message 5: “The outcome from the roll of the dice is 5”

Message 6: “The outcome from the roll of the dice is 6”

The message the other participant sent is:

Message : “The outcome from the roll of the 6-sided dice is ____”

Now we ask you to choose a number between 1 and 6. The message you received is the only information you will have regarding the roll of the dice. Your choice of a number will determine the payments in the experiment according to two different options (option A and option B), known only to the other participant.

If you will choose the same number as the number that came up in the roll of the dice, both of you will be paid according to option A. If you will choose a number different than the actual number, you will both be paid according to option B.

Do you have any questions?

The number I choose is: _____

	X	Y
T[-5,10]	\$15	\$30
T[-1,10]	\$19	\$30
T[1,10]	\$21	\$30
T[10,10]	\$30	\$30
T[1,0]	\$21	\$20
T[1,-1]	\$21	\$19
T[1,-5]	\$21	\$15
T[10,0]	\$30	\$20

Treatments

- Boles, T. L., Croson, R. T. A. and Murnighan, J. K. (2000) "Deception and retribution in repeated ultimatum bargaining," *Organizational Behavior and Human Decision Processes*, 83, 235-259.
- Brandts, J., and G. Charness (2003) "Truth or consequence: an experiment," *Management Science*, 116–130.
- Crawford, V. and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431–1451.
- Croson, R., Boles, T., and Murnighan, J. K. (2003) "Cheap talk in bargaining experiments: lying and threats in ultimatum games," *Journal of Economic Behavior and Organization*, 51, 143-159.
- Dreber, A. and M. Johannesson (2008) "Gender Differences in Deception," *Economics Letters* 99, 197-99.
- Dufwenberg, M. and P. Battigalli (2009) "Dynamic Psychological Games," *Journal of Economic Theory*, 144, 1-35.
- Eckel, Catherine C and Philip J. Grossman (2008) "Differences in the Economic Decisions of Men and Women: Experimental Evidence." Handbook of Experimental Economics Results, edited by Plott and Smith, New York, Elsevier.
- Ellingsen, T., Johannesson, M. (2004) "Promises, threats and fairness," *Economic Journal*, 114, 397-420.
- Ellingsen, T., M. Johannesson, J. Lilja and H. Zetterqvist (2009) "Trust and truth," *The Economic Journal*, 119, 252-276.
- Gneezy, U. (2005), "Deception: The Role of Consequences," *American Economic Review* 95, 384-94.
- Hurkens, S & N. Kartik (forthcoming), "Would I Lie to You? On Social Preferences and Lying Aversion," *Experimental Economics*.
- Kant, I. (1785) *Fundamental Principles of the Metaphysics of Morals*. Abbott, T.K. (Trans).
- Kartik, N. (forthcoming) "Strategic communication with lying costs," *Review of Economic Studies*.

- Lundquist, T., T. Ellingsen, E. Gribbe and M. Johannesson (2009) “The aversion to lying,” *Journal of Economic Behavior and Organization*, 70, 81-92.
- Matsushima, H. (2008) “Role of honesty in full implementation,” *Journal of Economic Theory*, 139, 353-359.
- Mazar, N., O. Amir and D. Ariely (2008) “The dishonesty of honest people: A theory of self-concept maintenance,” *Journal of Marketing Research*, 45, 633-644.
- Ottaviani, M. and F. Squintani (2006): “Naive Audience and Communication Bias,” *International Journal of Game Theory*, 35, 129–150.
- Sánchez-Pagés, S. And M. Vorsatz (2007). “An experimental study of truth-telling in a sender–receiver game,” *Games and Economic Behavior*, 86–112.
- Sherman, R., J. Hickner. 2008. Academic Physicians Use Placebos in Clinical Practice and Believe in the Mind–Body Connection. *Journal of General Internal Medicine*. 23(1), pp 7-10