Games and Economic Behavior 102 (2017) 433–444

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Incentives and cheating

Agne Kajackaite [a,*], Uri Gneezy [b,c]

[a] University of Southern California, 3620 S. Vermont Ave, Los Angeles, CA 90089-0253, United States
[b] University of California San Diego, 9500 Gilman Drive, La Jolla, CA 92093, United States
[c] University of Amsterdam, The Netherlands

ARTICLE INFO

Article history:
Received 22 July 2016
Available online 16 February 2017

JEL classification:
C91
D03
D63
D82

Keywords:
Cheating
Lying costs
Incentives
Laboratory experiment

ABSTRACT

We study how cheating behavior is affected by incentives. After replicating the finding in the cheating game literature that lying does not increase with incentives, we show that this insensitivity is not a characteristic of the intrinsic lying cost, but rather a result of concern about being exposed as a liar. In a modified "mind" game in which this concern is eliminated, we find that people lie more, and in particular lie more when the incentives to do so increase. Thus, our results show that for many participants, the decision to lie follows a simple cost–benefit analysis: they compare the intrinsic cost of lying with the incentives to lie; once the incentives are higher than the cost, they switch from telling the truth to lying.

1. Introduction

Much of the interaction in the marketplace is based on asymmetric information, which tempts some people to misreport to their advantage the private information they have. The economic consequences of cheating are large, and hence, understanding the factors that influence the decision to lie is important in order to understand many economic behaviors.

For some people, lying is intrinsically costly—they prefer an outcome if they reach it honestly over getting it by lying. Intrinsic costs of lying could result from a direct distaste of lying, or from reasons that are indirectly related to lying.[1] In principle, we can distinguish between three types of people according to their intrinsic cost of lying. Some people may be unwilling to tell a lie, regardless of their benefit from it ("ethical type"). Scholars such as St. Augustine (421) and Kant (1787) advocated such an uncompromising approach to lies. People who are not willing to lie could be described in our approach as having an infinite cost of lying. Other people may have a finite positive intrinsic cost of lying. These people will lie when the benefit of lying is higher than the associated cost ("finite positive cost type"); at the extreme are people with a zero cost of lying ("economic type").

* Corresponding author.
E-mail addresses: kajackai@usc.edu (A. Kajackaite), ugneezy@ucsd.edu (U. Gneezy).

[1] Such indirect preferences may relate to interdependent preferences in which the agent's utility depends on the utility of others—because of payoffs (Rabin, 1993; Fehr and Gächter, 2000) or because of guilt aversion (Dufwenberg and Gneezy, 2000; Charness and Dufwenberg, 2006; Battigalli and Dufwenberg, 2007; Battigalli et al., 2013). They may also relate to maintaining a positive social and/or self-image (Bénabou and Tirole, 2006; Andreoni and Bernheim, 2009), to the nature of communication (Lundquist et al., 2009), or to aversion to violating social norms (Elster, 1989; Bicchieri, 2006).

http://dx.doi.org/10.1016/j.geb.2017.01.015

In this paper we use an experiment to better understand how the size of the incentives to lie affects behavior. We do so by manipulating the incentives associated with lying and observing the effect on people's tendency to lie. The size of the incentive would not affect the ethical type who never lies or the economic type who always lies when doing so benefits them. By contrast, if the intrinsic cost of lying is positive but finite and the incentives are increased enough, some people will switch from telling the truth to lying.

The experimental economics literature offers two main approaches to studying lying behavior, and their results differ regarding the effect of incentives. The first line of research uses the "deception game" in which a sender has private information regarding payoffs associated with each of the choices that a receiver faces. The sender sends a message to the receiver that could be true or false, and the receiver decides whether to follow the message, which determines the payoffs for both players. In the deception game, the sender's action (message) may affect the receiver's beliefs; hence, this is a strategic decision. As a two-player game, the deception game also includes social preferences. The results are that senders are more likely to lie when the incentives to do so are increased (e.g., Gneezy, 2005; Sutter, 2009; Dreber and Johannesson, 2008; Erat and Gneezy, 2012).

The second approach to studying lying in the lab uses some form of a non-strategic "cheating game" in which the participant receives private information about the state of the world and reports it to the experimenter; this report, in turn, determines the participant's payoff. This type of game involves no strategic element. Furthermore, no social preferences are involved (assuming the participant does not take the experimenter's payoff into account). In contrast to the finding in the deception game, the literature on cheating games reports that higher incentives do not result in a higher fraction of lies (see Mazar et al., 2008; Fischbacher and Föllmi-Heusi, 2013, and the discussion in the literature review below).

In our experiments, we focus on the non-strategic games, i.e. cheating games. We first replicate the result of the cheating games by testing a game in which a participant rolls a six-sided die in private and then reports the result to the experimenter. If the participant reports a 5, she receives \$X (\$X = \$1, \$5, \$20 and \$50, between participants), and zero otherwise. Cheating is detected by comparing the reported distribution of numbers with the theoretically predicted one. In line with the findings reported above, we find no increase in cheating as the incentives are raised.

Why is behavior affected by incentives in deception games but not in cheating games? One possibility is that in the deception game, lying is an option suggested in the rules of the game where the sender is asked to decide whether to send a true or false message. The consequence of each message is well defined in terms of the strategic outcome of the game. This is not the case in the cheating game, in which participants are asked to report the actual outcome without being told about the option to lie. Because the lying option is not explicitly mentioned in the rules of a cheating game, participants might be afraid of possible consequences of a lie. In addition, participants in the cheating game might question whether their die roll is really private, or whether the experimenter has a way to find out the outcome. When the incentives are raised, concerns may increase regarding the possible consequences of being exposed when lying (e.g., Ekman, 1985, 1988; Vrij, 2008).

To test the effect of this potential concern, we introduced a "mind game" in which we ask participants to think about a number in private, then roll the die in private, and report whether the number that came up is the same as the one they thought of (see Jiang, 2013; Shalvi and De Dreu, 2014, and Potters and Stoop, 2016, for a similar manipulation). If participants report it was the same number, they receive \$X as before, and zero otherwise. As predicted, the results of this "mind game" show an overall increase in participants' tendency to lie for every level of X. This result indicates that some participants in the cheating game chose not to lie because of the risk of being exposed, and not because of the intrinsic cost of lying. In addition, and in contrast to the cheating game, in the mind game, the fraction of participants who lied increases with X. That is, more participants lied when the incentives were increased. The mind game presents an extreme case of zero probability of being exposed. While it is hard to imagine such situations in daily life, the extreme case is helpful in testing and understanding the intrinsic cost of lying and its interaction with incentives.

Our results provide an explanation for findings in the literature on cheating games, and show that in a mind game the tendency to lie increases with incentives, indicating that some of our participants have positive and finite intrinsic costs of lying. In particular, the evidence suggests that lying is a "normal good" in which people compare the intrinsic cost and benefit of the lie, and when the benefit from lying is higher than the intrinsic cost of lying, they lie.

2. Literature review

As discussed in the introduction, the evidence from cheating games suggests that increasing the financial incentives does not increase lying. Mazar et al. (2008) asked participants to solve a set of 20 matrices in private and report their success to the experimenter later on. The participants received payment based on a piece rate per correctly solved matrix. Payments were increased, between participants, from \$.50 to \$2 (from a maximal payoff of \$10 to \$40). The results of this treatment were compared with a baseline test of ability, and show participants lie, but that there is no difference between the two levels of incentives.

Even stronger than this finding, Mazar et al. (2008) report the results of additional experiments in which they manipulated the amount of payment to each participant per correctly solved matrix (\$.10, \$.50, \$2.50, and \$5). They find, compared with the control, a "limited dishonesty" in the \$.10 and \$.50 treatments, but no lying in the \$2.50 and \$5 conditions. Mazar et al. (2008) conclude behavior is consistent with a convex cost function: the cost of lying increases faster than the benefits, such that when the incentives to cheat are increased, people are not more likely to lie.

Fischbacher and Föllmi-Heusi (2013) introduce the die-roll method, in which participants roll a die in private and then report the outcome to the experimenter. The participants' payoff depends on their report: they receive 1, 2, 3, 4, 5, 0 Swiss Francs for reporting 1, 2, 3, 4, 5, 6, respectively. To test the effect of incentives, in one of the control treatments, all payments were multiplied by a factor of 3 (from a maximal payoff of 5 Swiss Francs to 15 Swiss Francs). Fischbacher and Föllmi-Heusi (2013) find that the fraction of participants who lie does not change significantly when the stakes are tripled. The authors conclude, "Either the effect of increasing the benefits of lying is counterbalanced by the negative effect of increased costs of lying when a person earns more through deception, or the concept of lying is not directly related to stakes at all."

Abeler et al. (2016) provides a meta-analysis of the die-roll game. With respect to incentives, they find (based to a large degree on the results we report in this paper), that "[S]ubjects continue to refrain from lying maximally when stakes are increased. Indeed, we find almost no increase in the average report when stakes are increased 500-fold." (Abeler et al., 2016, p. 8).

Related to this, Weisel and Shalvi (2015) introduce a cheating game with an element of cooperation—two paired players have to report the same outcome sequentially in order to receive a payoff, where the payoff equals the number reported. Weisel and Shalvi (2015) show that a substantial fraction of participants collaborates and lies. However, doubling the incentives does not affect the lying behavior.

By contrast, in the experiment reported in Gibson et al. (2013), each participant acted as a CEO and had to report earnings per share for the previous quarter. The earnings had two states, and reporting the false state resulted in the higher payoff for the participant. The incentive to lie was varied such that the payoff for the wrong report was the same (1.5 Swiss Francs), but the payoff for a truthful message differed over the five treatments from 0.3 to 1.5 Swiss Francs. Gibson et al. (2013) find that lying increases from 18% up to 79%, when it becomes more beneficial. However, the design is very different from the simple abstract studies of Fischbacher and Föllmi-Heusi (2013) and Mazar et al. (2008) because of the framing used, the use of a within-subject design (Charness et al., 2012), and telling participants that reporting wrong earnings was legal.

Together, the experimental evidence from cheating games suggests there is no effect from the size of financial incentives on lying, unless the game is modified such that it is more similar to the deception game. In particular, Gibson et al. (2013) tell the participants that lying is legal—in the deception game, lying is also an option the rules of the game suggest. By contrast, if one uses a standard cheating game, the fraction of people who lie is not increasing with the incentive to lie and might even be decreasing.

3. Experimental design, hypotheses and procedure

We concentrate on the effect of changing the incentives and measuring the effect of these changes on lying behavior. In the following, we describe the theoretical framework, treatment variations and hypotheses.

3.1. Theoretical framework

To illustrate the decision situation an agent faces in cheating and mind games, we derive a simple utility function. Consider a player who observes the state of nature t and then reports the state. She can either report the true state t or report a false state t'. If she reports the true state, her monetary payoff is m_t and if she lies it is $m_{t'}$. Thus, the monetary benefit of a lie is $m_{t'} - m_t$.

Denote the intrinsic cost of lying as C_i. If a player takes into account only the monetary benefit of a lie and has no intrinsic cost of lying ($C_i = 0$), she would lie whenever $m_{t'} - m_t > 0$ and be indifferent between lying and telling the truth if $m_{t'} - m_t = 0$. If the player is an ethical type with $C_i = \infty$, she would not take into account how large $m_{t'} - m_t$ is and would always report truthfully. Finally, there is a third type of player who has a finite and positive cost of a lie: $0 < C_i < \infty$. She will lie whenever $m_{t'} - m_t > C_i$. With respect to the functional form of the intrinsic lying cost, we make a simplifying assumption that the intrinsic lying cost is fixed for each individual, i.e. the size of the lie, which we determine by $t' - t$ and $m_{t'} - m_t$ does not affect the lying cost (see Gneezy et al., 2016).

Additionally, the player experiences disutility if being exposed when lying, which she believes will happen with probability p_i. We assume that the player's "fear" of getting caught is a function $f(m_{t'}, m_t, p_i)$ which is increasing in $m_{t'} - m_t$ and p_i. We make an assumption that the fear function is convex in stakes, i.e. we assume that marginal fear with respect to the reported payoff increases in the reported payoff. Furthermore players are heterogeneous with respect to the affinity experiencing this fear. We denote this affinity by parameter γ with $\gamma \geq 0$.

Put together, a player will lie whenever the utility below is positive:

$$m_{t'} - m_t - C_i - \gamma f(m_{t'}, m_t, p_i) > 0.$$

This utility function represents our idea that cheating behavior depends on p_i. As we argued above, players in the cheating game have a higher p_i than players in the mind game, and this will motivate our hypotheses on the differences between cheating and mind games.

Table 1
The structure of the cheating and mind games.

Basic Cheating Game	**Basic Mind Game**
• Roll a die in private • Report the number • If report "5" receive \$X; otherwise receive \$0	• Think of a number • Roll a die in private • Report "Yes" if the number is the same as you thought of; report "No" otherwise • If report "Yes" receive \$X; otherwise receive \$0

Note: The structure of the cheating game is described on the left side, and the structure of the mind game on the right side. In both experiments X = \$1, \$5, \$20 and \$50 between subject. 75 observations per cell (600 in total).

3.2. Basic cheating game

The cheating game we use is a variation of Fischbacher and Föllmi-Heusi's (2013) experiment described above, in which a participant rolls a six-sided die in private, and then reports the outcome to the experimenter. In the game, we told the participant that if she would report a 5, she would receive \$X; otherwise, she would earn \$0. In a between-participants design, we vary X such that \$X = \$1, \$5, \$ 20, and \$50. The procedure of the cheating game is described on the left side of Table 1.

In such a cheating game, participants might be concerned about being exposed when lying (denoted by $f(m_{t'}, m_t, p_i)$ in the utility function) and this concern might increase with incentives. Participating in the lab experiment such as the cheating game for \$1, \$5 or even \$20 is a normal procedure for many students. Participating in this experiment and receiving \$50 for reporting the lucky outcome is unusual, and might make participants question what the consequences will be if he or she lies.

The concern about being exposed when lying could come from, for example, standard preferences regarding penalties. This concern may positively depend on the size of the monetary gain from lying. According to the social psychology literature, and in line with intuition, stakes are predicted to increase the fear of getting caught when lying (e.g., Ekman, 1985, 1988; Vrij, 2008). If the concern about being exposed when lying is high enough, the incentive effect in the cheating game will not be positive.

The hypothesis we test in this game is in line with the findings from cheating games discussed above. We expect the increase in incentives to lie (increasing $m_{t'} - m_t$) will not lead to more lying, since moving from \$1 to \$50 will also increase the fear of being caught ($f(m_{t'}, m_t, p_i)$):

Hypothesis 1. Increasing Incentives in the cheating game will not increase lying.

3.3. Extended cheating game: cheating game with a tedious task

There is a connection between cheating and punishment (e.g., Brandts and Charness, 2003; Houser et al., 2012). In a sense, people might use being treated badly as a justification to lie. Having such an excuse to cheat may increase cheating. To test this prediction, we gave our participants a tedious retyping task before they took part in the same cheating game as in the first treatment. We gave them piles of 200 questionnaire forms and asked them to type in the data in an excel sheet. We asked them to work for 15 minutes. We explicitly told the participants they would not receive payment for the task, hence creating a feeling of unfairness. The payoffs and the rules of the cheating game were the same as in the first treatment.

Hypothesis 2. Being treated unfairly before playing the cheating game will increase cheating.

3.4. Basic mind game

In the mind game treatments, the participant is first asked to think about a number between 1 and 6 and remember the number (without reporting it to us). She then rolls a six-sided die. If the die outcome is the same as the number she thought of, the participant is asked to report "Yes" and otherwise to report "No." If the participant reports "Yes," she receives \$X as above. If she reports "No," she receives \$0. The procedure of the mind game is described on the right side of Table 1.

In the mind game participants know their lie will not be exposed. This reduction in the fear of being exposed suggests two hypotheses. First, with respect to the overall lying, participants will lie more in the mind game than in the cheating game, because the fear element disappears:

Hypothesis 3. Participants will lie more in the mind game than in the cheating game.

Second, as discussed above, because the fear about being exposed is eliminated, we also anticipate a positive effect of incentives:

Hypothesis 4. Increasing incentives in the mind game will increase lying.

Table 2
Experimental summary.

Type of game	Payoffs	N			
		\$1	\$5	\$20	\$50
Basic cheating game	Receive \$X if report 5; \$0 otherwise.	75	75	75	75
Extended cheating game	Tedious task before cheating game. Receive \$X if report 5; \$0 otherwise	75	75	75	75
Basic mind game	Receive \$X if report yes; \$0 otherwise	75	75	75	75
Extended mind game	Receive \$X if report yes, and other player receives \$0; \$0 otherwise and other player receives \$50	75	75	75	75

3.5. Extended mind game: two-players mind game

As discussed above, one major difference between the deception game and the cheating and mind games is the effect the decision to lie has on another participant's payoffs. In the next treatment, we test the effect of social preferences by introducing a second player whose payoff depends on the participant's decision to lie. In this two-player version of the mind game, if the participant reports "Yes," she receives \$X and the other participant assigned to her receives zero; if she reports "No," she receives zero and the other participant receives \$X. Having the second player in the game may affect cheating because reporting "No" results in envy for the first player, and reporting "Yes" leads to compassion, where envy is supposed to be stronger than compassion (e.g., Fehr and Schmidt, 1999):

Hypothesis 5. Participants will lie more in a two-player mind game than in the basic mind game.

This treatment also addresses the concern that participants do not cheat because they care about the payoff of the experimenter. Because the experimenter pays no matter what the participant reports, such a concern is not applicable here. Also, note that we do not encourage lying in neither of the treatments and do not put the option to lie in the instructions of the experiment.

3.6. Experimental procedure

We conducted the experiments in January–March 2015 at the Rady Behavioral Lab, UC San Diego, and recruited participants through Sona Systems. A total of 1200 participants (48.3% female) took part in the experiment, with 75 observations per treatment (with a total of 16 treatments differing by the type of a game ($\times 4$) and the size of incentives ($\times 4$)), and none of them participated in more than one session. We set N such that the probability of correctly rejecting the null hypothesis of no lying when the null hypothesis is false, is at least 80%. More specifically, we set the number of observations such that if we observe 30% reporting of a "5" (or "Yes"), the binomial test has a power of more than 80% to reject the null hypothesis of no lying. Using $p = 0.167$ (expected) and $p = 0.3$ (observed), $N(75)$ leads us to power of 84.34%.

Participants played the cheating/mind game at the end of an experimental session. At the beginning of our experiment, participants received written instructions (full instructions are reported in Appendix A) and were allowed to ask questions privately. Then each participant rolled the die in private and reported the outcome on a sheet of paper. After completing the task, participants were asked to complete a post-experiment questionnaire that included questions on gender, age, and field of study. At the end, participants privately received their payoffs in cash and left the laboratory. Table 2 summarizes the experimental treatments and procedure.

4. Results

In what follows, we examine the incentive effects on lying in the cheating and mind games. The left bars of Fig. 1 display the fraction of participants who reported a die outcome of 5 in the cheating game. When the incentive is \$1, 25% of participants claim to have rolled a 5. Slightly more participants report a 5 when the incentive is \$5 or \$20 with 29% and 33% of positive reports, respectively. This positive trend in the fraction of participants who cheat is in line with the hypothesis that cheating will increase with incentives. However, raising the incentives to \$50 reverses this trend. When the incentive is increased to \$50, the reporting of a 5 declines to 23%. The pairwise between-treatments differences are not significant in a Fisher test ($p > 0.1$, two-sided). The overall trend is also not significant—when testing for an order (both ascending or descending) in the four-incentives conditions, we find no effect of stakes on the fraction of lies ($p > 0.1$, Jonckheere–Terpstra). The trend is also not significant when considering only the \$1, \$5 and \$20 conditions, in which the incentives have a slight positive effect ($p > 0.1$, Jonckheere–Terpstra).

The expected number of participants who rolled a 5 is 16.67%; hence, in expectation, any value above it is based on lies. We run a binomial test and find participants lie significantly for \$1, \$5, and \$20 with $p = 0.061$, $p = 0.008$, and $p < 0.001$, respectively (all tests reported in the paper are two-tailed). In the \$50 treatment lying becomes insignificant with $p = 0.164$. This result is in line with the concern about being exposed increasing with incentives and the fear function being convex with respect to stakes such that participants do not lie when the stakes are increased by a factor of 50. In other words, it is

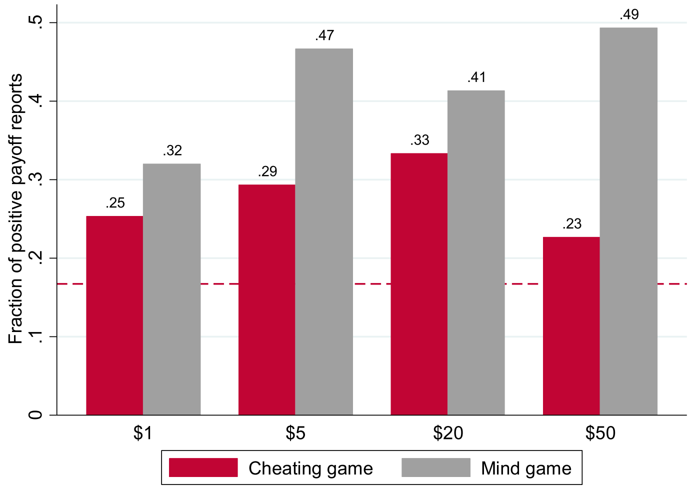

Note: The left columns represent levels of cheating (above the expected 16.7% marked by the dashed line) in the cheating game by incentives. The right columns represent these numbers for the mind game.

Fig. 1. Lying in the basic cheating and mind games.

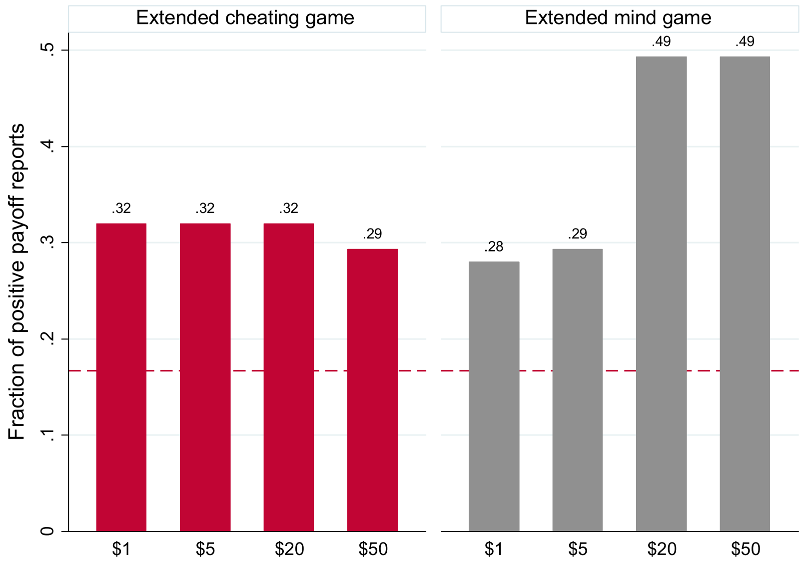

Note: The left columns represent levels of cheating (above the expected 16.7% marked by the dashed line) in the extended cheating game by incentives. The right columns represent these numbers for the extended mind game.

Fig. 2. Lying in the extended cheating and mind games.

consistent with the assumption that in the cheating game fear is increasing in $m_{t'} - m_t$ and in particular the marginal fear with respect to the reported payoff increases in the reported payoff.

In the cheating game with a tedious task treatment, we find that 32% of participants report a 5 when the incentive is $1, $5, and $20, and 29% report a 5 for $50 dollars (see Fig. 2). As in the main cheating game, the pairwise between-treatments differences are not significant in a Fisher test ($p > 0.1$). When testing for an ascending order in a Jonckheere–Terpstra test, we find no significant effect of stakes on the fraction of lies ($p > 0.1$). Finally, we compare the theoretical distribution with reported outcomes, and find lying is significant for each size of the incentive in a binomial test on a 1% level.

We run a Probit regression to check for the robustness of the incentive effects in the cheating games (Table 3 lists the results). In the regression, we also include the dummy *Real-effort stage* to control whether lying is affected in the cheating game with a tedious task. As can be seen from the first model (columns 1 and 1m), we find the same results in the regressions as in the non-parametric tests. The insignificant coefficients of the *$5, $20 and $50 dummies* in the model show that in the cheating game the size of incentives has no effect on the proxy of lying—reporting a positive payoff. The regression also shows lying in the cheating game does not change if we introduce the tedious task stage.

Table 3
Probit regression models for cheating and mind games.

	(1)	(1m)	(2)	(2m)
Dependent variable	Reporting positive payoff	Reporting positive payoff	Reporting positive payoff	Reporting positive payoff
Type of the game	Cheating	Cheating	Mind	Mind
Model	Probit	Probit (ME)	Probit	Probit (ME)
\$5	0.049	0.017	0.250*	0.098*
	[0.153]	[0.053]	[0.149]	[0.059]
\$20	0.124	0.043	0.421***	0.165***
	[0.153]	[0.054]	[0.148]	[0.058]
\$50	−0.072	−0.025	0.519***	0.203***
	[0.156]	[0.053]	[0.148]	[0.058]
Real-effort stage	0.145	0.050		
	[0.110]	[0.038]		
Second player			−0.079	−0.031
			[0.105]	[0.041]
Male	0.288***	0.099***	0.152	0.059
	[0.111]	[0.038]	[0.105]	[0.040]
Constant	−0.786***		−0.580***	
	[0.140]		[0.134]	
Observations	599	599	599	599
Pseudo R-squared	0.013	0.013	0.020	0.020

Note: The models are Probit specifications with the fraction of positive payoff reports ("5" or "Yes") as the dependent variable. Columns 1m and 2m report marginal effects from Probit models. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. \$5, \$20, \$50, Male, Real-effort stage, and Second player are dichotomous variables equal to 1 if the participant is in the \$5 treatment, \$20 treatment, \$50 treatment, the participant is male, the participant worked on a real-effort task, and we introduced a second player, respectively. The reference group in Model 1 consists of participants in the cheating game. The reference group in Models 2 consists of participants in the mind game.

Based on this regression result, and the fact that all pairwise non-parametric comparisons between the basic cheating game and the cheating game with a tedious task are also not significant ($p > 0.1$, Fisher test), we pool the data from the two treatments and use it as the pooled cheating-game data in the following analyses; to be conservative, we also report the analyses of the separate treatments in the text. The above analysis reveals our data is in line with Hypothesis 1. In particular, we find the following:

Result 1. Increasing incentives in the cheating game does not increase lying.

The results reject Hypothesis 2, since working on a tedious task before playing the cheating game did not increase cheating.

Moving from the cheating game to the mind game, we conjectured two effects: first, overall lying will increase; second, the fraction of participants who lie will depend positively on incentives.

The right bars in Fig. 1 present the results from the basic mind game in terms of the fraction of participants who reported "Yes." In line with Hypothesis 3, we observe that more participants lie in the basic mind game than in the basic cheating game. In the basic mind game, 32%, 47%, 41%, and 49% of participants report a positive outcome for \$1, \$5, \$20, and \$50, respectively, whereas in the basic cheating game, the positive reporting rate varied between only 23% and 33%. For the stakes of \$5 and \$50, the between-treatments difference is significant in a Fisher test with $p = 0.043$ and 0.001, respectively. Other pairwise comparisons are not significant ($p > 0.1$). When we compare the fractions of positive reports in the pooled cheating-game data (basic cheating game and extended cheating game) and the basic mind game treatment, we observe the same effects—for each size of the stake, participants lie more in the basic mind game than in the pooled cheating game data, and for stakes of \$5 and \$50, the difference is significant in a Fisher test, with $p = 0.027$ and 0.001, respectively. In addition, in the basic mind game, lying is significant for each size of the incentive in a binomial test ($p < 0.001$).

In the two-player mind game, participants also lie more than in the cheating games. Here, 28% of participants report "Yes" for \$1, and 29% do so for \$5, and the fraction of positive reports amounts to 49% for both \$20 and \$50 (see Fig. 2). In this game, lying is significant in a binomial test for each size of the stakes on the 1% level. When we compare the lying in the two-player mind game with lying in the pooled-cheating game data, we find the fractions of positive reports are higher in the two-player mind game, when the stakes are \$20 and \$50, with $p = 0.020$ and $p = 0.001$ (Fisher test), respectively, and are not different for \$1 and \$5. When comparing the lying in the basic cheating game (not pooled) with the two-player mind game, all effects hold—we find the fractions of positive reports are higher in the two-player mind game than in the basic cheating game, when the stakes are \$20 and \$50, with $p = 0.068$ and $p = 0.001$ (Fisher test), respectively, and are not different for \$1 and \$5 ($p > 0.1$).

We run a Probit regression to find further support for lying being higher in mind games than cheating games. Table 4 reports the regression results. We include the dummy *Mind game* to measure whether participants lie more in the mind game. As can be seen from column 1m, the marginal effect of dummy *Mind game* is positive, significant, and large in size. In particular, we find, on average, participants lie 14.7 percentage points more in the mind game than in the cheating game.

Table 4
Probit regression model for overall lying.

	(1)	(1m)
Dependent variable	Reporting positive payoff	Reporting positive payoff
Model	Probit	Probit(ME)
Mind game	0.401***	0.147***
	[0.107]	[0.039]
$5	0.159	0.060
	[0.108]	[0.041]
$20	0.278***	0.105***
	[0.107]	[0.041]
$50	0.238**	0.090**
	[0.107]	[0.041]
Real-effort stage	0.137	0.051
	[0.109]	[0.041]
Second player	−0.072	−0.026
	[0.104]	[0.038]
Male	0.218***	0.080***
	[0.076]	[0.028]
Constant	−0.889***	
	[0.114]	
Observations	1198	1198
Pseudo R-squared	0.022	0.022

Note: The model is a Probit specification with the fraction of positive payoff reports as the dependent variable. Column 1m reports marginal effects from the Probit model. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. $5, $20, $50, Male, Real effort stage, Second player, and Mind game are dichotomous variables equal to 1 if the participant is in the $5 treatment, $20 treatment, $50 treatment, the participant is male, the participant worked on a real-effort task, we introduced a second player, and the participant played the mind game, respectively. The reference group consists of participants in the cheating games.

Considering the utility function we introduced above, the difference between the two games is in the probability of being caught p_i. The data show that lowering p_i by moving from the cheating game to the mind game makes participants lie significantly more. In particular, our data do not reject Hypothesis 3:

Result 2. A larger fraction of participants lie in the mind game than in the cheating game.

Note that this result is not surprising. While the main focus of this paper is the effect of incentives, the result shows that, as expected, the probability of being caught affects the tendency to lie.

In the regression models, we also control for gender effects. We use regression models in Table 3 to measure these gender effects. We find that men lie significantly more in the cheating game—the coefficient of *Male* is significant, positive, and large (see column 1m, Table 3). This gender effect, however, disappears when we move to the mind game—the insignificant coefficient of *Male* in column 2m (Table 3) shows men do not lie more than women in the mind game.[2]

We did not anticipate that women and men will lie to the same degree in the mind game, but women will lie less in the cheating game. An ad hoc explanation of this result is related to gender differences in risk-aversion literature, where women are shown to be more risk averse (e.g., see Charness and Gneezy, 2012). A possible interpretation of the results is that in the cheating game, the risk of getting caught keeps women from lying, whereas the risk affects men to a lower extent. As soon as this risk is eliminated, women and men lie to the same degree.

After we show that a larger fraction of participants lie in the mind game than in the cheating game, the second conjecture we have with respect to mind games is that lying will increase with incentives in the mind game (Hypothesis 4); the data supports this hypothesis. When the gain from reporting "Yes" in the basic mind game is $1, 32% of participants report it. This fraction increases to 47% ($5), 41% ($20), and 49% ($50). The pairwise between-treatments comparisons show a marginally significant difference between lying for $1 and $5, and a significant difference between $1 and $50, with $p = 0.094$ and $p = 0.046$, respectively. Other pairwise comparisons are not significant ($p > 0.1$). When we test for an ascending order in a Jonckheere–Terpstra test over the four conditions, we find a significant increasing effect of stakes on the fraction of lies, with $p = 0.034$. That is, participants lie more when the stakes are raised in the mind game.

[2] In general, gender differences in lying are still relatively unexplored. Dreber and Johannesson (2008) find that in situations where the sender can benefit at the expenses of the receiver by lying ("selfish black lies"), men are more likely to lie than women. Erat and Gneezy (2012) replicated this result, but found that in situations when the sender loses but the receiver gains in the case of a lie, women were more likely to lie. Conrads and Lotz (2015) find that women are more responsive to changes in communication channels (see Croson and Gneezy, 2009). Muehlheusser et al. (2015) used a similar die-rolling paradigm as in Fischbacher and Föllmi-Heusi (2013), dividing participants up into groups of two. They find no gender differences in cheating, but more cheating in male groups and mixed groups than in female groups.

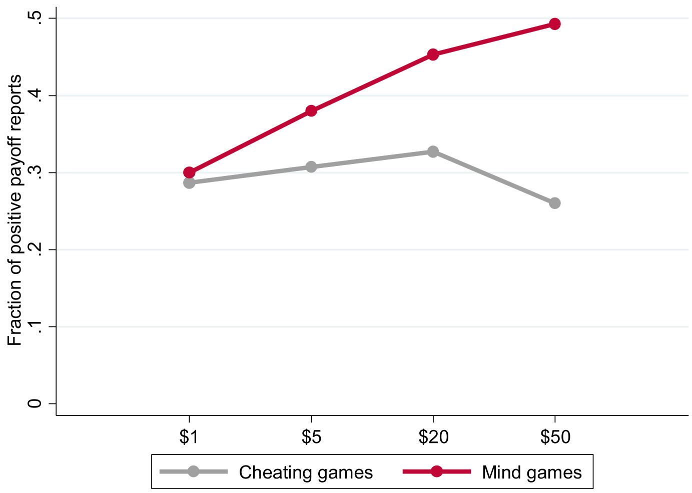

Note: The lines represent levels of cheating (above the expected 16.7%) in the cheating and mind games by incentives.

Fig. 3. Lying in the cheating and mind games groups.

In the two-player mind game, lying also increases with stakes. The fraction of participants reporting "Yes" goes from 28% and 29% for \$1 and \$5, respectively, up to 49% in the \$20 and \$50 conditions (see Fig. 2). The pairwise between-treatments comparisons show a significant difference between lying for \$1 and \$20, \$1 and \$50, \$5 and \$20, and \$5 and \$50, with $p = 0.012$, $p = 0.012$, $p = 0.019$, and $p = 0.019$, respectively; other pairwise comparisons are not significant ($p > 0.1$). When we test for an ascending order in a Jonckheere–Terpstra test, we find a highly significant effect of stakes on the fraction of lies, with $p < 0.001$.

In the robustness checks using a Probit regression (see Table 3), as can be seen from Model 2 (columns 2 and 2m), we find a significant incentives effect in the mind games. The marginal effects of *\$5, \$20 and \$50* dummies are significant and increasing with the size of the incentive (column 2m). The regression model also shows lying in the mind game does not change if we introduce the second player. Using pairwise non-parametric comparisons between the basic mind game and the two-player mind game, we find between-treatments differences are not significant for \$1, \$20, and \$50 ($p > 0.1$, Fisher test), but the difference is significant for \$5, with $p = 0.043$. Based on the insignificant *Second player* coefficient in the regressions and the fact that most of the non-parametric test results above are insignificant, we pool the data from both mind games for our tests, as well as report the non-pooled data tests. Based on the comparisons, we also reject Hypothesis 5, as we do not find that in the two-player mind game participants were more likely to lie than in the basic mind game.

Fig. 3 presents the pooled results for the cheating games and mind games. Fig. 3 shows a clear large incentive effect on lying in the mind games but not in the cheating games. The non-parametric tests results from the pooled data, which are the same as in the separate data analyses, support this effect. In particular, a highly significant positive incentive effect on lying exists in the mind games ($p < 0.001$, Jonckheere–Terpstra), whereas no such effect exists in the cheating games ($p = 0.719$, Jonckheere–Terpstra).[3]

The above analysis reveals that our data cannot reject Hypothesis 4, presumably because in the mind game the fear is close to zero (or equal to zero) and is only marginally affected by $m_{t'} - m_t$. In particular, we find the following:

Result 3. Increasing incentives in the mind game increases lying.

5. Conclusion

Although the economics literature has extensively studied lying, many questions regarding individual lying behavior remain unanswered, in particular, the effect of incentives on cheating. The evidence reported in the literature show that people do not cheat more when the incentives to do so are increased. This evidence suggests that people have either zero lying costs and will always lie when doing so benefits them, or infinite costs and will never lie. However, when we asked people in the audience during a seminar about their type, all but one person (out of over 100) self-identify as being someone who has a finite lying cost and would lie if the benefit outweighs the cost.

[3] An additional reason for a higher fraction of lying in the mind game than in the cheating game might be self-deception: it might be easier for the player to convince herself that she got the "right" number in the mind game than in the cheating game. However, self-deception cannot explain the incentives effect in the mind game.

A possible explanation for the observation that people do not lie more when the incentives are increased is based on negative emotions, such as guilt, that affect the identity of the person who lies (Mazar et al., 2008; Fischbacher and Föllmi-Heusi, 2013). However, we show that when we modify the standard cheating game to eliminate the concern about being exposed as a liar, participants do react to the incentives to lie. Since the negative emotions should not be different between the cheating game and the mind game, these emotions cannot explain the results.

The difference in behavior between the games indicates that the finding in the literature regarding the intrinsic lying cost picked up an additional element in the utility of the participants: Increasing the incentives to lie also raised the participants' concern about being exposed. Whereas our results show that incentives increase cheating in the mind game where these concerns about being exposed are eliminated, they also show that this effect is limited, because at least half of our participants chose not to lie. The large fraction of people who did not lie could be caused by our incentives not being large enough, by intrinsic lying costs being very high, by an additional concern that we did not identify, or because many people are just not willing to lie for money.

It is interesting to compare the above result to that of the telephone experiment reported in Abeler et al. (2014), in which participants were paid based on self-reported coin flips over the phone. The design was such that it "was obvious that the true outcome was only known to the participants, as they tossed the coin privately at home" (p. 97). In contrast with our mind game, Abeler et al. (2014) found very little cheating in their game. Understanding why cheating behavior is different across these two games and environments (students in the lab versus a representative sample of the German population in a phone experiment) would be important for further understanding the structure of lying costs.

Another interesting observation from our experiment is that the $50 treatments seem to be the outlier that breaks monotonicity in the cheating game. The results of the mind game reject the moral sentiments explanation used in the literature: It is not simply that large incentives make participants averse to cheating. We argue, instead, that large incentive may make cheating look criminal and hence is more deterring, for example because participants perceive $50 to be high enough to justify detection of liars. However, we have only four incentives level tested in this paper. What would happen, for example, if we offered participants $500 if they cheat? Future research can answer this and other related questions regarding the shape of the "demand function" for cheating.

Together, the results show that ethical behavior associated with the intrinsic cost of lying follows standard economic reasoning in our experiments. When benefit from lying increases, more people choose to lie. These findings have important implications for constructing policies to reduce cheating. When the interaction is closer to the cheating game one, in which peoples' concern regarding being exposed increases with incentives, it might be better to make the incentives associated with the decision to cheat high. On the other hand, when the interaction is closer to the mind game one, in which people are less concerned about the chance of being exposed, it might make sense to reduce the incentives to cheat.

Acknowledgments

We thank Lea Cassar, Gary Charness, Patrick Kampkötter, Erin Krupka, Silvia Saccardo, Anne Schielke, Marina Schröder, Shaul Shalvi, Dirk Sliwka, Joel Sobel, Charlie Sprenger, Marcin Waligora, and seminars and conferences participants for valuable comments and suggestions.

Appendix A. Experimental instructions

The instructions include only the treatments with a $1 incentive.

A.1. Cheating game

Thank you for taking part in our experiment. This experiment is part of a study we are conducting, and your help is highly appreciated.

We have given you a die, and ask you to roll it in private after you read the instructions.

We will then ask you to write at the bottom of this form the number that you rolled. Your payment will be $0 if you write "1", "2", "3", "4," or "6," and it will be $1 if you write "5." That is, if you write "5," you will be paid $1; if you write a different number, you will not be paid. The payoffs are as follows:

Number thrown	1	2	3	4	5	6
Payoff	$0	$0	$0	$0	$1	$0

After you have rolled the die, please write down the number you rolled. Then come to the other room of the lab to fill out a post-questionnaire and receive your payment. Please bring this form with you.

Do you have any questions?

Please roll the die in private now and report the number that came up:_____

Thank you for participating!

A.2. Cheating game with a tedious task

Part 1

Thank you for taking part in our experiment. This experiment is part of a study we are conducting, and your help is highly appreciated. This experiment consists of two parts. After you have finished the first part, you will receive the instructions for the second part.

Part 1: We gave you a packet with 200 copies of questionnaire forms. Participants in a previous experiment filled out these questionnaires. Your task is to copy the data into the excel sheet on your computer screen.

First, please write down your participant number on the upper left side. Also, please write down what time you started working on the task and what time you finish. We ask you to work on the task for 15 min.

We ask you to fill in six columns: “participant number,” “gender,” “ethnicity,” “age,” “field of study,” and “year.” If you cannot read a questionnaire form, please just continue to the next one.

At the end of the 15 minutes, please raise your hand and a research assistant will come over to you to save the excel file and give you the instructions for the second part of the experiment.

Do you have any questions?

Thank you!

Part 2

These are the instructions for the second part of the experiment.

We have given you a die, and ask you to roll it in private after you read the instructions.

We will then ask you to write at the bottom of this form the number that you rolled. Your payment will be $0 if you write “1”, “2”, “3”, “4,” or “6,” and it will be $1 if you will write “5.” That is, if you write “5,” you will be paid $1; if you write a different number, you will not be paid. The payoffs are as follows:

Number thrown	1	2	3	4	5	6
Payoff	$0	$0	$0	$0	$1	$0

After you have rolled the die, please write down the number you rolled. Then come to the other room of the lab to fill out a post-questionnaire and receive your payment. Please bring this form with you.

Do you have any questions?

Please roll the die in private now and report the number that came up:_____

Thank you for participating!

A.3. Mind game

Thank you for taking part in our experiment. This experiment is part of a study we are conducting, and your help is highly appreciated.

We have given you a die, and ask you to roll it in private after you read the instructions.

Now think of a number from 1 to 6 and remember it. Then, at the bottom of this form, please write the number that you rolled at the bottom of this form.

If the number you rolled is the number you thought of, please mark “Yes” at the bottom of the form, and your payment will be $1. If the number you rolled is not the number you thought of, please mark “No” at the bottom of the form, and your payment will be $0.

After rolling the die, writing down the number you rolled, and marking “Yes” or “No,” please come to the other room of the lab to fill out a post-questionnaire and receive your payment. Please bring this form with you.

Do you have any questions?

Please roll the die in private now and report the number that came up:_____

Please indicate whether the number you rolled is the number that you thought of:

☐ Yes

☐ No

Thank you for participating!

A.4. Two-player mind game

Thank you for taking part in our experiment. This experiment is part of a study we are conducting, and your help is highly appreciated. In the experiment, you will be matched with another participant.

We have given you a die, and ask you to roll it in private after you read the instructions.

We ask you now to think of a number from 1 to 6 and remember it. Then, at the bottom of this form, please write the number you rolled.

If the number you rolled is the number you thought of, please mark “Yes” at the bottom of the form. Your payment will be $1, and the person you are matched with will be paid $0. If the number you rolled is not the number you thought of, please mark “No” at the bottom of the form. Your payoff will be $0, and the person you are matched with will be paid $1.

After rolling the die, writing down the number you rolled, and marking "Yes" or "No," please come to the other room of the lab to fill out a post-questionnaire and receive your payment. Please bring this form with you.

Do you have any questions?

Please roll the die in private now and report the number that came up:_____

Please indicate whether the number you rolled is the number that you thought about:

☐ Yes

☐ No

Thank you for participating!

References

Abeler, J., Becker, A., Falk, A., 2014. Representative evidence on lying costs. J. Public Econ. 113, 96–104.
Abeler, J., Nosenzo, D., Raymond, C., 2016. Preferences for Truth-Telling. IZA Discussion Paper 10188.
Andreoni, J., Bernheim, B.D., 2009. Social image and the 50–50 norm: a theoretical and experimental analysis of audience effects. Econometrica 77 (5), 1607–1636.
St. Augustine John, 421, Enchiridion on Faith, Hope and Love, Regnery Publishing, Inc., Washington, DC, 1996.
Battigalli, P., Dufwenberg, M., 2007. Guilt in games. Amer. Econ. Rev. 97 (2), 170–176.
Battigalli, P., Charness, G., Dufwenberg, M., 2013. Deception: the role of guilt. J. Econ. Behav. Organ. 93 (C), 227–232.
Bénabou, R., Tirole, J., 2006. Incentives and prosocial behavior. Amer. Econ. Rev. 96 (5), 1652–1678.
Bicchieri, C., 2006. The Grammar of Society: The Nature and Dynamics of Social Norms. Cambridge University Press.
Brandts, J., Charness, G., 2003. Truth or consequences: an experiment. Manage. Sci. 49, 116–130.
Charness, G., Dufwenberg, M., 2006. Promises & partnership. Econometrica 74, 1579–1601.
Charness, G., Gneezy, U., 2012. Strong evidence for gender differences in risk taking. J. Econ. Behav. Organ. 83 (1), 50–58.
Charness, G., Gneezy, U., Kuhn, M., 2012. Experimental method: between-subject and within-subject design. J. Econ. Behav. Organ. 81, 1–8.
Conrads, J., Lotz, S., 2015. The effect of communication channels on dishonest behavior. J. Behav. Exp. Econ. 58, 88–93.
Croson, R., Gneezy, U., 2009. Gender differences in preferences. J. Econ. Lit. 47, 448–474.
Dreber, A., Johannesson, M., 2008. Gender differences in deception. Econ. Letters 99 (1), 197–199.
Dufwenberg, M., Gneezy, U., 2000. Measuring beliefs in an experimental lost wallet game. Games Econ. Behav. 30, 163–182.
Ekman, P., 1985. Telling Lies. Clues to Deceit in the Marketplace, Politics and Marriage. Norton, New York.
Ekman, P., 1988. Lying and nonverbal behavior: theoretical issues and new findings. J. Nonverbal Behav. 12 (3), 163–175.
Elster, J., 1989. Social norms and economic theory. J. Econ. Perspect. 3 (4), 99–117.
Erat, S., Gneezy, U., 2012. White lies. Manage. Sci. 58 (4), 723–733.
Fehr, E., Schmidt, K.M., 1999. A theory of fairness, competition and cooperation. Quart. J. Econ. 114 (3), 817–868.
Fehr, E., Gächter, S., 2000. Fairness and retaliation: the economics of reciprocity. J. Econ. Perspect. 14 (3), 159–181.
Fischbacher, U., Föllmi-Heusi, F., 2013. Lies in disguise – an experimental study on cheating. J. Eur. Econ. Assoc. 11 (3), 525–547.
Gibson, R., Tanner, C., Wagner, A.F., 2013. Preferences for truthfulness: heterogeneity among and within individuals. Amer. Econ. Rev. 103 (1), 532–548.
Gneezy, U., 2005. Deception: the role of consequences. Amer. Econ. Rev. 95 (1), 384–394.
Gneezy, U., Kajackaite, A., Sobel, J., 2016. Lying aversion and the size of the lie. Working paper.
Houser, D., Vetter, S., Winter, J., 2012. Fairness and cheating. Europ. Econ. Rev. 56 (8), 1645–1655.
Jiang, T., 2013. Cheating in the mind games: the subtlety of rules matters. J. Econ. Behav. Organ. 93, 328–336.
Kant, I., 1787. On a supposed right to lie from benevolent motives. In: Beck, Lewis W. (Ed.), The Critique of Practical Reason and Other Writings in Moral Philosophy. University of Chicago Press, Chicago, 1949, pp. 364–350.
Lundquist, T., Ellingsen, T., Gribbe, E., Johannesson, M., 2009. The aversion to lying. J. Econ. Behav. Organ. 70, 81–92.
Mazar, N., Amir, O., Ariely, D., 2008. The dishonesty of honest people: a theory of self-concept maintenance. J. Marketing Res. 45 (6), 633–644.
Muehlheusser, G., Roider, A., Wallmeier, N., 2015. Gender differences in honesty: groups versus individuals. Econ. Letters 128, 25–29.
Potters, J., Stoop, J., 2016. Do cheaters in the lab also cheat in the field? Europ. Econ. Rev., 26–33.
Rabin, M., 1993. Incorporating fairness into game theory and economics. Amer. Econ. Rev. 83, 1281–1302.
Shalvi, S., De Dreu, C.K.W., 2014. Oxytocin promotes group serving dishonesty. Proc. Natl. Acad. Sci. USA 111, 5503–5507.
Sutter, M., 2009. Deception through telling the truth?! Experimental evidence from individuals and teams. Econ. J. 119, 47–60.
Vrij, A., 2008. Detecting Lies and Deceit: Pitfalls and Opportunities. John Wiley, Chichester.
Weisel, O., Shalvi, S., 2015. The collaborative roots of corruption. Proc. Natl. Acad. Sci. USA 112, 10651–10656.