

Jiang, Ting

Working Paper

The Mind Game: Invisible Cheating and Inferable Intentions

LICOS Discussion Paper, No. 309

Provided in Cooperation with:

LICOS Centre for Institutions and Economic Performance, KU Leuven

Suggested Citation: Jiang, Ting (2012) : The Mind Game: Invisible Cheating and Inferable Intentions, LICOS Discussion Paper, No. 309, Katholieke Universiteit Leuven, LICOS Centre for Institutions and Economic Performance, Leuven

This Version is available at:

<https://hdl.handle.net/10419/75021>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



LICOS Centre for Institutions and Economic Performance

Centre of Excellence

LICOS Discussion Paper Series

Discussion Paper 309/2012

The Mind Game : Invisible Cheating and Inferable Intentions

Ting Jiang



Katholieke Universiteit Leuven

LICOS Centre for Institutions and Economic
Performance
Waaistraat 6 – mailbox 3511
3000 Leuven

BELGIUM

TEL: +32-(0)16 32 65 98

FAX: +32-(0)16 32 65 99

<http://www.econ.kuleuven.be/licos>

The Mind Game: Invisible Cheating and Inferable Intentions

Ting Jiang*

LICOS, KULeuven; CentER, Tilburg University

Abstract

This paper exploits a novel cheating game - the “Mind Game” - to show how a subtle variation in the rules of the game affects cheating. In both variants of the game, cheating is invisible because subjects make their choices purely in their minds. The only difference rests on the order of the steps in which subjects should play the game. I find that subjects cheat significantly less when they have to disregard the prescribed order of steps in order to cheat, compare to when they can purely lie about the choice made in the mind. Since subjects play the game in private with a double-blind payment procedure, I conjecture that subjects only cheat to the extent that their moral self-image is still preserved. But the moral image can be preserved if they can deceive themselves into believing that at least their intentions are good, and not if cheating requires an act that reveals the intent to cheat. This study thus suggests a potential role of intent inference in deterring cheating.

Keywords: cheating, intent, self-deception, moral self-image, the mind game

JEL codes: C91, D63, H26, K42

*I wish to thank Jan Boone, Eric van Damme, Jan Potters for their invaluable comments and suggestions. I also want to thank Juan-Juan Cai, Koen Deconinck, Ernesto Reubens, Sigrid Suetens and Yang Zhao for their helpful comments. Special thanks go to Jan Willem Lindemans for his helpful comments on the design of the mind game as well as the earlier drafts of the paper. I also want to thank participants in TIBER Symposium 2011, ESA Luxembourg 2011 and Nordic Conference 2011 for their feedbacks on the paper presentation. I am grateful to the Netherlands Organization for Scientific Research (NWO) for its generous financial supports with the Mozaiek grant. Contact: ting.jiang@econ.kuleuven.be.

1. Introduction

Most people manage to morally self-regulate, albeit imperfectly. Mazar, Amir and Ariely (2008)'s seminal paper shows that even honest people cheat, though only a little bit, so as to maintain a good self-image. Needless to say, the economic loss of small cheating by the majority, though negligible at the individual level, can be daunting when aggregated (Ariely, 2008; DePaulo et al., 1996). More insights on how to deter “honest” people from cheating can yield substantial economic and social benefits.

The theory of self-image maintenance highlights the relevance of “self” being an observer of the cheating act and deepens our understanding of the peculiar cheating patterns of the honest that standard economic theory fails to explain (Mazar et al., 2008)¹. People do not cheat more when they are less likely to get caught by others since they can be caught by themselves (see also, Fischbacher & Heusi, 2008). They are also unlikely to cheat more when the external reward of cheating is higher if cheating more would then lead to a higher internal moral cost (such as more feelings of guilt) than the extra monetary benefit (see also, Lundquist et al., 2009; Shalvi, Handgraaf & De Dreu, 2011)². But, people seem to cheat more when intricate psychological mechanisms such as self-justification and self-deception enable them to think well of themselves while cheating (e.g., Chance et al., 2011; Shu, Gino & Bazerman, 2011)³.

To illustrate this with a hypothetical example, imagine an employee who casually took some pens home from the office. When questioned about why he took the pens home, he might reply: “I did not even notice that I took the pens.” “They are almost worth nothing. How can I possibly benefit from it in any significant way?” “Besides, it is only fair since I sometimes work at home.”

Probably, the employee embraced these excuses himself for preserving a moral self-image, even if no one noticed the missing pens or questioned him. In daily life, one can notice that people often safeguard their moral self-image with the help of self-deceptive excuses (such as blaming it falsely on ignorance and errors, or tricking oneself to believe that one is only claiming one's just or fair

¹ Gneezy (2005) points out that the extent of lying also depends on the harm on others.

² Note that this result is more applicable to cheating behavior in non-strategic settings. Another study by Abler, Becker & Falk (2012) also found little cheating outside the lab in a non-strategic context.

³ Recent economic theories not only theoretically investigated the rationale of maintaining a positive self-image, but also adopted self-deception as a crucial factor of self-image maintenance (e.g., Benabou and Tirole 2002).

share)⁴. To deter cheating by those who care about self-image, it partly boils down to the question of how to make the context of cheating less susceptible to self-deceptive excuses.

This paper exploits two novel variants of a cheating game, which I called the “Mind Game”, to create two different cheating contexts. The Mind Game is a game in which subjects make a choice purely in their mind. Since subjects get paid based on their self-reported choice, they can cheat by lying about the choice⁵. The game consists of three steps: “choose a side”, “throw a die” and “report the side chosen”. In both variants of the game subjects first choose which side of a six-sided die will count for their earnings: the side facing up or the side facing down. The only difference between the two variants rests on the prescribed order of the last two steps. In the “throw-first” variant, after choosing a side in their mind, subjects *first throw* the die, and only after they see the outcome they report the side chosen, “Up” or “Down”. In the “report-first” variant, the order prescribes subjects to *first report* their chosen side and then throw the die. In both variants, by repeating the game twenty times, individuals’ likelihood of cheating can be inferred from the proportion of “lucky” choices. However, the nature of cheating is different between the two variants. In the “throw-first” variant, subjects can legitimately know the lucky side when reporting. In contrast, in the “report-first” variant, subjects cannot know the lucky side unless postponing the step of reporting. This study shows how such subtle variation in the rules of the game affects the extent of cheating.

Although the overall results suggest a consistent pattern of small cheating and the absence of full extent cheaters, the two cheating contexts exhibit distinctive cheating patterns. Above all, I find that subjects in the “report-first” treatment cheat significantly less. Since subjects play the game in private with a double-blind payment procedure, I conjecture that this is so because the moral self-image that people are said to care about is based on the recognition that at least one’s intentions are good. And it is the intent or motive behind the act that determines the moral judgment of the act, thus leaving room for self-deceptive and self-justifying excuses. Although the cheating act in both contexts is undeniable to oneself, the intent to cheat is more salient in the “report-first” treatment and less in the other.

⁴ For instance, House et al. (2011) found that subjects cheat more if they were treated unfairly before.

⁵ In this study, cheating is defined as subjects’ reporting untruthfully about the side chosen in their minds which results in a higher payoff than that of the honest side.

In the “throw-first” treatment, the intent to cheat is unclear. Subjects who repeatedly report the “lucky” choice can maintain a positive self-image by arguing that they had forgotten to choose a side in mind or which side they chose, or that they had accidentally reported the lucky side. In the “report-first” treatment, the act of disregarding the order of play cannot be easily accounted for by errors or ignorance but is inferably due to the bad intent to cheat. Subjects thus have less room for excuses that can “save their intentions”. In order to use the excuse of “unintentional cheating”, a subject would need to have accidentally disregarded the order of playing the game in addition to the excuses mentioned above. Thus, the intent to cheat of subjects who cheated is almost undeniable. The necessary preparatory act for cheating makes it harder for subjects to deceive themselves about their intent to cheat for preserving a moral self-image while cheating.

To illustrate this using the earlier example of stealing pens, suppose that, instead of casually taking the pens when others are present, the employee would have waited until the office was empty and carefully looked around before reaching for the pens. Then, he would probably fail to safeguard his moral self-image with the same excuses because of his almost pronounced intent to steal, revealed to him by his own “sneakiness”.

That intentions matter for moral judgment is not new. In the interpersonal cheating context, intention has already been shown to be crucial in determining whether the deceived is willing to forgive (e.g., Schweitzer et al. 2006), and people tend to refrain from intentional cheating (Bandura, 1990; Bandura et al., 1996). In the common law tradition, *mens rea* (“the guilty mind”) is a crucial element for a certain act to formally constitute a crime⁶. From an evolutionary perspective, self-deception has even been argued to serve an adaptive purpose of hiding cheaters’ bad intents (von Hippel and Trivers, 2011).

So far, experimental studies, mostly in strategic settings, have demonstrated that intentions matter for reaching cooperative outcomes⁷. Experiments on individual cheating, however, mainly

⁶ The full sentence goes, “the act is not culpable unless the mind is guilty”, which is translated from Latin, “*actus non facit reum nisi mens sit rea*”.

⁷ It has been shown that asymmetric information about intention leads to less altruism, and communication of intention can promote cooperation by evoking guilt aversion (e.g., Charness & Dufwenberg, 2006). Falk, Fehr and Fischbacher (2002), for instance, showed that the intentions of an agent are important for the extent of reward and punishment of that agent. See also Falk, Fehr & Fischbacher (2008) for behavioral evidence on how attribution of fairness intentions matters.

focused on the detection of the cheating act but not the intent to cheat, most likely because intentions are invisible. The fact that intentions are not directly observable is probably why bad intents can be easily covered up by self-deceptive excuses and why cheating is so pervasive despite of people being lying-averse. It is thus all the more important to understand how intent detection affects cheating.

This paper provides some supportive experimental evidence that in a cheating context in which subjects are more directly confronted with their own intentions to cheat, they cheat less.

Moreover, the distinctive intricate cheating patterns across contexts are likely to result from different cheating strategies for maintaining a moral self-image, triggered by the ease of intent inference.

The mind game devised in this paper is useful for studying subtle cheating. Compared to the latest cheating game in experimental economics in which subjects self-report the die-throw outcome in private (Heusi & Fischbacher, 2008), there are two main improvements: first, since the act takes place only in the mind, the mind game creates a cheating context with zero possibility of being caught without any additional set-up needed to ensure secrecy. Note that only in the “report-first” treatment, subjects can still fear the exposure of the preparatory act of changing the prescribed order. In both treatments, subjects need not fear the exposure of their cheating, even if they sit in public or even under camera surveillance. Second, less data are needed for inferring the cheating probability since the mind game is a binomial process similar to that of a coin flip. Though developed independently, it resembles, to a certain extent, the game implemented in a neuroscientific experiment by Greene and Paxton (2009) in which subjects self-report the prediction of a coin flip outcome. Their aim, however, was to study the psychological mechanisms behind honesty and hence the details of the game as well as the treatments differed (see Section 2 for more details).







As a policy implication, my results suggest that rules should be designed so that there is little room for unintentional cheating and few opportunities for self-deceptive excuses. Such a mechanism is almost costless relative to external enforcement mechanisms, since it relies on people’s moral self-regulation. Sometimes even small changes in the rules of the game, like changes in the order of doing things, are enough to render intentions more inferable and thus to deter cheating.

2. Experimental Design

2.1 The Mind Game

The essence of the mind game lies in making a choice in mind before knowing the outcome of chance. An abstract form of the game consists of the following three stages: one makes a choice in mind without revealing it; nature moves by chance; one self-reports the initial choice made. The mind game in this study is an individual cheating game.⁸ It uses a normal six-sided die.⁹ Correspondingly, subjects first have to choose purely in the mind which side of the die will count for their earnings: the side facing up or the side facing down. Then they throw the die, and only after subjects see the outcome, they have to write down the side initially chosen, “Up” (“U”) or “Down” (“D”). For each outcome of the die, the earning points can be either high or low corresponding to the side chosen. For instance, if the outcome of the throw is “1” and the subject reports, truthfully or not, that he had “D” in mind, he receives “6” in payoff. The exact earnings for all six outcomes corresponding to the two mind choices can be found in Table 1.

Table 1: Earning Points in the Mind Game

						
U	1	2	3	4	5	6
D	6	5	4	3	2	1

Note: The number of dots on the opposite sites of a six-sided die always adds up to 7.

Subjects can cheat by reporting a side that will give them more points: “4”, “5”, “6” instead of “1”, “2”, “3”. Because the choices were made purely in subjects’ minds, cheating is invisible to others. By repeating the game multiple times, cheating can even be inferred at the individual level if the proportion of “lucky” choices is improbably high, relative to the expected theoretical level of 50%. The mind game is different from the Greene and Paxton (2009)’s coin flip paradigm in the sense that it combines a binary choice in mind with various die outcomes which enables a wide range of investigations on intricate cheating patterns (see Section 4 for more details). In particular,

⁸ Note that the mind game can also be run to examine cheating in strategic settings (see Jiang, 2012 for an application of the mind game in a bribery setting.)

⁹ When playing with physical dice or a coin, a die has a small advantage over a coin, namely, the outcome is less subject to the skill of flipping a coin randomly, which might add noise.

the mind game features an embedded variation of the cheating gain. Since the number of dots on the opposite sides of a six-sided die always add up to 7, the cheating gain ranges from five points of difference with the die outcomes of “6” and “1”, three points with “5” and “2” and one point with “4” and “3”. The non-intrusive implementation of varying the cheating gain provides an embedded within-subject treatment of the cheating payoff size and serves to study more subtle and automatic cheating “strategies”.

2.2 Opacity and Cheating

I examine how cheating is affected by a subtle variation of the rules of the game by switching the order of the last two steps “throw the die” and “report the side”. In the “Opaque” (“throw-first”) treatment, the step “throw the die” is placed before “report the side” to enable subjects to cheat purely in mind. In the “Transparent” (“report-first”) treatment, subjects have to *report* the chosen side on paper immediately after choosing the side in their minds. Only then can they “*throw* the die” *and see* the die outcome (see Table 2 for the steps).

Table 2: Steps of the two mind game variants

	Opaque (“throw-first”)	Transparent (“report-first”)
Step 1:	choose a side in mind	choose a side in mind
Step 2:	throw the die	fill in the chosen side on paper
Step 3:	fill in the chosen side on paper	throw the die

It is reasonable to expect that subjects hold the belief that the experimenter has no access to their minds in both treatments. Though the act of disregarding the order of the steps in Transparent is visible for oneself, it is impossible for the experimenter to find out if subjects follow these steps or not since the experimenter stays in another room most of the time. Even if the experimenter could have all of a sudden walked towards a certain subject, all three steps would have been finished within a few seconds before the experimenter arrived.

The key difference rests on subjects’ own awareness of their intent to cheat. In Transparent, because the explicit act of disregarding the prescribed order cannot be easily accounted for by

errors or ignorance but the purpose to cheat¹⁰, one is forced to be aware of the cheating intent conditional on the cheating act. The intent to cheat is less obvious in Opaque, since subjects can legitimately see the die outcome before reporting the side. Cheating only requires an internal twist of the mind without giving explicit cues for inferring the intent to cheat.

Strictly speaking, the possibility exists in Transparent to be caught on the intent to cheat also by the experimenter, though not in Opaque. As an example, suppose that the experimenter would stand behind a subject throughout the experiment in Transparent. Although the experimenter can never verify the subject's claim that he or she was honest about the sides, the experimenter can catch the subject in not following the order and infer the intent to cheat irrespective of the hidden cheating act. All in all, the probability of cheating's being intentional is higher in Transparent for any given probability of cheating. The probability of catching one's own cheating act is bigger than zero, but the same across treatments. The difference lies in the opaqueness of the cheating contexts for intent inference.

2.3 Experimental Procedure

The experiment was conducted in the CentER lab in September 2010. There were six sessions and 43 subjects in total participated¹¹. A session lasted 30 minutes on average including 20 rounds of the mind game as well as a short questionnaire. On average, subjects earned 6 Euro. A general introduction to the experiment and instructions were read aloud to subjects, as well as a short description of how they were going to be paid (see Appendix I for the instructions)¹². Moreover, they were asked to come up with a personal password for picking up their payment in the

¹⁰ Note that the act of disregarding the order itself is not sufficient to constitute cheating. The definition of cheating used in this study distinguishes the act of cheating from the mere act of breaking a given rule since breaking a rule does not always imply that subjects have cheated. For instance, if a subject fills out "Up" for all the 20 rounds at the beginning of the experiment without following the rule of doing it round by round, he or she is not perceived as a cheater in this study since he or she cannot possibly benefit from the act of breaking this rule.

¹¹ Four out of forty-three subjects cannot be included in the analysis: two self-reported confusion about the game and expressed disagreement with the answer to the understanding test in the post-experimental survey; another two came up with a same password which made it impossible to match their choices of "Up" and "Down" to the die outcomes since the password was the only source of information for matching the die outcomes and the choices of sides.

¹² Subjects were told that the experiment was about "individuals' successive responses towards outcomes that are randomly generated in a lottery setting".

secretaries' office one week later. In this way, anonymity toward the experimenter was assured. The experimenter also has no possibility to make any statistical inference about the honesty of the subjects when subjects submit their outcome forms since only the sides "U" or "D" were written on the outcome form, but not the die outcomes.

The instructions were read to them in the waiting hall followed by a test of understanding. Afterwards, they were led into the computer room and sat in front of separated computers. After filling in their personal passwords on the welcoming page for picking up their payment, they proceeded to the actual experiment and played the game at their own speed. The exact procedure of how subjects implement the three steps in Opaque was as follows (see Appendix III for the screen shots): on the first screen, subjects read, "If you have chosen the side ("U" or "D") in your mind, please click on the button below". After clicking on the button "throw the die", subjects entered the second screen in which the die outcome was revealed graphically. Then subjects were supposed to fill in the side on the outcome form before they clicked on the second button "next round". In Transparent, whereas all else remained the same, subjects were asked to fill in the side already on the first screen, "If you have chosen the side ("U" or "D") in your mind, please fill in the side on the outcome form and click on the button below:". For both treatments, the die outcomes were randomly generated in each round for each subject and stored on the web server.

3. Main Result

Result 1: *Subjects in Opaque cheat significantly more than those in Transparent based on both the foresight and the earnings levels.*

There are a number of ways to infer the probability of cheating in the mind game from the experimenter's perspective, irrespective of how subjects catch themselves. The most simple and straightforward approach is to examine whether the sides subjects choose yield earnings of "4", "5", "6" improbably more than "1", "2", "3", as if they had "foresight" of the die outcomes before

the die is thrown. First, let f_{ij} be the indicator of earnings higher than three for individual i and round j ; let individual “foresight” F_i denote the average of f_{ij} over the 20 rounds of individual i :

$$F_i = \frac{\sum_{j=1}^{20} f_{ij}}{20}, \text{ where } f_{ij} = 1 \text{ if earning} > 3; 0 \text{ otherwise.}$$

Likewise, the treatment foresight is the average foresight over all rounds of the treatment in question. As the die is fair, the expected foresight is 0.5. Thus, inference of cheating can be drawn when the foresight significantly deviates from the theoretical level of 0.5. Opaque clearly exhibits improbably high foresight ($F_O = 0.64$) over all rounds (one-tailed binomial test, $p < 0.001$) and over individuals (Wilcoxon signed-ranked test, $p = 0.001$). The significance levels are lower in Transparent ($F_T = 0.55$) (one-tailed binomial test, $p = 0.041$; Wilcoxon signed-ranked test, $p = 0.107$).

Above all, the average foresight levels, shown in Table 3, show that subjects in Opaque cheat significantly more than those in Transparent based on the average individual foresights [Mann-Whitney-Wilcoxon (MWW) test, $p < 0.05$]. At the aggregate level, Opaque also exhibits significantly more cheating than Transparent (Pearson’s chi-square test, $p < 0.01$).

Table 3: Treatment Effects on Cheating Levels

	Opaque	Transparent	Test of difference (MWW)
Average Individual Foresight	0.64	0.55	$p=0.032$
Average Individual Earning	0.64	0.54	$p=0.055$

A second measurement is related to total earnings. Since subjects are randomly exposed to cheating opportunities of earning five, three or one points corresponding to different die outcomes, the actual realization of die throws over the 20 rounds is not the same across subjects. To capture the proportion of the actual earning relative to the maximal possible earning with cheating given the individual’s draw of die throws, the earning measurement is defined as the aggregate of the actual earnings per round e_{ij} normalized over the difference of the aggregate maximum earnings ($F_i = 1$) and minimal earnings ($F_i = 0$):

$$E_i = \frac{\sum_{j=1}^{20} e_{ij} - \sum_{j=1}^{20} \underline{e}_{ij}}{\sum_{j=1}^{20} \bar{e}_{ij} - \sum_{j=1}^{20} \underline{e}_{ij}}, \text{ where } \underline{e}_{ij} = (e_{ij} | f_{ij} = 0); \bar{e}_{ij} = (e_{ij} | f_{ij} = 1).$$

E_i is between 0 and 1. Without cheating, E_i is also expected to be 0.5. Again, Opaque exhibits improbably high earning levels ($E_o = 0.64$) based on the individual earnings (Wilcoxon signed-ranked test, $p = 0.001$). Transparent, however, does not exhibit clear cheating ($E_T = 0.54$, Wilcoxon signed-ranked test, $p = 0.139$). As shown in Table 3, subjects in Opaque also cheat more than those in Transparent according to average individual earnings, with the P-value slightly above 0.05¹³. The difference in significance levels according to the two measurements is potentially due to specific disguising strategies in place: if subjects try to disguise earning by cheating only for low gains, cheating would still be captured by the foresight measure, and vice versa. It is thus reassuring to use both measurements.

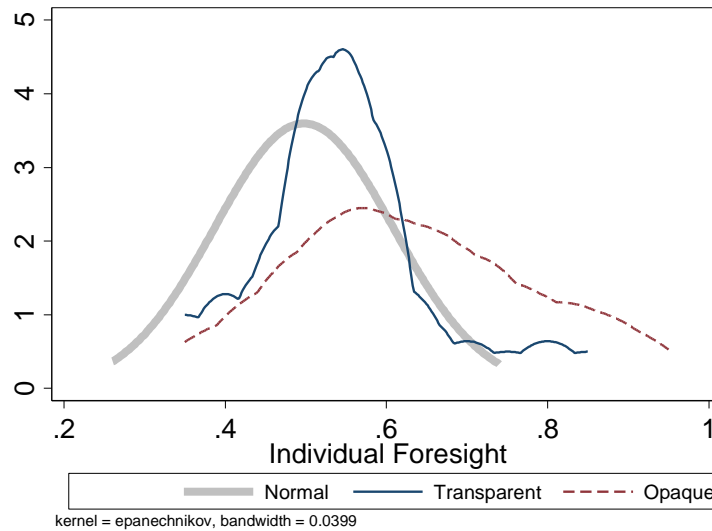
More support for the different cheating levels is also found based on the comparison of the number of “cheaters”, defined as subjects exhibiting improbably high individual foresight ($F \geq 0.7$). Opaque features a significantly higher proportion of cheaters [33% vs. 11%, Fisher’s exact test (FE), $p = 0.1$], as well as a significantly smaller proportion of honest subjects (52% vs. 89%, FE, $p = 0.018$)¹⁴. The results are robust using the earnings measurement.

Note that, in the light of the “small cheating” found in previous studies, if subjects sometimes cheat so little that it is not even inferable statistically, it is then still possible that there are as many people who cheat in Transparent except that they cheat much less, such as for one or two rounds. As shown by the Kernel density estimates of the individual foresight (see Fig. 1), Transparent has a denser distribution between 0.5 and 0.6 relative to the asymptotically estimated normal distribution (Kurtosis test, $p = 0.16$), pointing to a plausible small cheating threshold around 0.6. The distribution of Opaque, however, is shifted rightward with a much thicker tail above 0.6 which implies that at least a substantial proportion of subjects cheated more than just a little bit.

¹³ The prevalence of cheating in Treatment O is also confirmed by the distribution of individual foresight in Treatment O which differs significantly from the truthful distribution (Kolmogorov-Smirnov test, $p=0.02$), while the difference is not significant in Treatment T, which again implies insignificant cheating.

¹⁴ 27 subjects with the lowest foresight levels, are classified as “honest” being the largest group of subjects that exhibit no significant evidence of cheating at the group level (average foresight = 0.52, one-tailed binomial test, $p>0.1$). But for those who have less than the expected 0.5 foresight level by chance, they can still cheat to reach the 0.6 foresight level while belonging to the honest group.

Figure 1: Treatment Effects on Distributions (Kernel density estimate)



All in all, though it is not certain that less subjects cheated in the transparent cheating context, the results strongly suggest that subjects cheat to a much smaller extent. And if there are cheaters in “disguise” in Transparent, they must have cheated so little that it is almost negligible. Overall, no one cheats 100% according to both measurements. The foresight and earning levels are about 28% above the expected level in Opaque and 10% in Transparent. Most intriguingly, the cheating level in Opaque (28%) corresponds to the elicited 20 to 31 percent lying frequency in daily life found by DePaolo et al. (1996). It raises an interesting question of whether the process of lying in daily life is similar to that in Opaque in the sense of only requiring an internal twist of the mind. In the next section, the analysis will zoom in on the subtle cheating behavior in relation to opaqueness.

4. Other Results

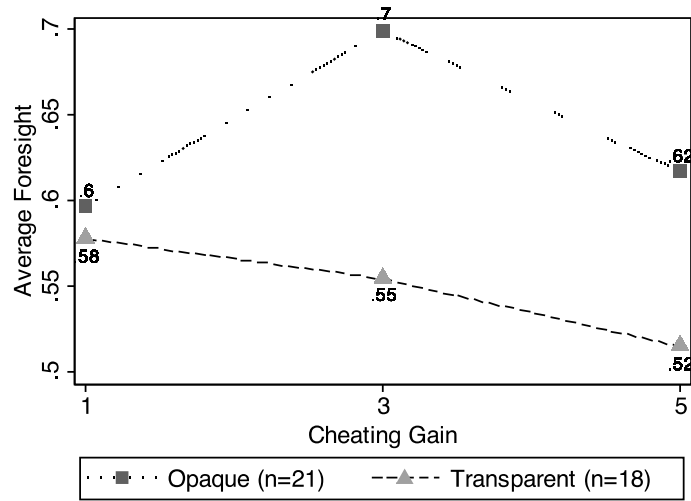
4.1 Opaqueness and Cheating Patterns

Result 2: *While significant cheating can only be found for small gains in Transparent, it is found for all levels of gains in Opaque.*

Previous studies have shown slightly different results on the exact cheating patterns corresponding to different cheating gains. While Fischbacher and Heusi (2008) found that subjects avoid cheating for the highest gain, Mazar et al. (2008) showed that subjects cheat for the two lower gains, but not for the two higher gains. A recent study by Shalvi et al. (2011) shows that subjects avoid both major and minor lies. Here the patterns seem to differ systematically in the two contexts.

As depicted in Figure 2, the average foresights seem to differ both across and within treatments, corresponding to cheating gains of 5, 3 and 1.

Figure 2: Average Foresight and Cheating Gains



In Transparent, significant deviation from the theoretical level of 0.5 is only the case for the average foresight of small gains when die outcome are 3 and 4 ($F_T^{3,4} = 0.58$, one-tailed binomial test, $p = 0.06$). This result confirms that subjects in Transparent not only cheat less frequently, but also cheat only when the reward is small. In Opaque, cheating is found at all levels of gains since the average foresights are all significantly higher than 0.5. Moreover, the average foresight seems to be at its peak for the medium gains ($F_T^{2,5} = 0.7$). Nevertheless, the difference in foresights is only significant between the medium and small gains (Chi-square, $p = 0.067$), but not between the medium and the high gains (Chi-square, $p = 0.147$). This indicates that subjects do not systematically avoid cheating for certain level of gains in Opaque.

Though these results may appear to be puzzling at first, it seems to fit in nicely with the treatment difference on intent inference. The high moral cost incurred in Transparent, because of the saliency of the intent to cheat, can be so high that the moral image is at risk if subjects cheat more than just a little bit. In Opaque, subjects cheat more randomly because there is probably no need for them to be selective upon certain cheating gains if self-deceptive excuses already reconcile their cheating and the moral image. In some sense, a subject cannot deceive himself that he does not intend to cheat while still strategically choosing which levels of gains to cheat on, unless if he uses heuristics to “choose”. Arguably, the opaqueness explanation can also shed light on the different extents of cheating found in previous studies. For instance, it is plausible that the small cheating pattern found in Mazar et al. (2008) is related to its rather transparent cheating context: since subjects grade their own tests for getting paid, the intent to cheat can be easily inferred if they cheat a lot. In Fischbacher and Heusi (2008), however, since subjects only roll the die once, the earnings vary substantially among subjects. The bad intent of cheating for undeserved benefits can be potentially covered up by the self-deceptive excuse that the rule itself is unfair.

Result 3: *Systematic patterns of “U” and “D” are only found in Transparent.*

Intriguingly, some subjects use specific patterns in choosing the side, such as choosing the same side throughout the 20 rounds or changing sides systematically (see Table 4 for the exact patterns).

Table 4: Patterns of Chosen Sides

	Sequence of the sides chosen (20 rounds)
1	UUUUUUUUUUUUUUUUUUUU
2	UUUUUUUUUUUDDDDDDDDDD
3	UDUDUDUDUDUDUDUDUDUD

Notes: Patterns are defined as such that they do not allow for more than 2 rounds of cheating. The pattern “UUUUUDDDDDDUDUDUDUDUD” used by one of the subjects was not included here as it allows for at least three times the possibility to switch patterns to one’s advantage.

Moreover, all subjects who used such patterns are found in Transparent (28% vs. 0%, FE test, $p = 0.01$)¹⁵. In a way, subjects in Transparent might have anticipated some high moral cost and

¹⁵ There was one subject in Opaque who reported using a pattern of “choosing the other side whenever a side fails to yield a lucky outcome”. But his actual pattern of sides did not correspond to the self-reported pattern.

immediately decided to commit to being honest in the beginning of the game. Future studies with more subjects can be conducted to confirm the persistent treatment difference of this pattern and investigate the underlying reasons.

Result 4: *Subjects in Opaque cheat on all die outcomes except “6”, while subjects in Transparent only cheat on die outcome “3”.*

For the rest who do not use patterns and choose more randomly “U” and “D”, results for cheating on specific outcomes can be obtained if the proportion of “U” chosen is unusually high or low ¹⁶. This is based on the fact that without foresight on die outcomes, the propensity of subjects’ choosing “U” or “D” should be the same for each outcome. I use the proportion of upsides chosen when foresight is zero as the benchmark to approximate the honest proportion of “U” chosen. On average, subjects in the honest rounds of Opaque chose “U” 51% of the time, whereas subjects in the Transparent chose “U” 55% of the time excluding the pattern users. There are two main results (see Table 5): first, subjects in Opaque cheat on all die outcomes except on “6”. Second, subjects in Transparent only cheat on die outcome “3”.

Table 5: Cheating and Die Outcomes

Side (Proportion of “U”)	Average (f=0)	1 point		3 points		5 points		All
		3	4	2	5	1	6	
Opaque	0.51 (n=151)	0.4**	0.59*	0.38**	0.78***	0.33***	0.55	0.51
Transparent	0.547 (n=117)	0.4**	0.64	0.44	0.55	0.59	0.56	0.58

Notes: *, **, *** indicates significance at the 90%, 95%, and 99% level respectively, one-sided binomial test. The 10 percent significance level is used in order to capture small cheating.

Both of the results hint at the unexpected asymmetry of the cheating propensity within each level of gain. Why would subjects only avoid cheating on “6” but not on “1” since cheating on “1” also yields 5 points? Perhaps “6” is simply too salient to cheat on, which also suggest that the cheating strategies that subjects used may be based on heuristics rather than on systematic deliberations.

¹⁶ Note that the foresight measurement is no longer appropriate here due to the inequality of upsides and downsides chosen. For instance, if there is a stronger tendency of choosing “U” irrespective of cheating, which is the case in Transparent, the foresight level will be upward biased when outcomes are 4, 5 or 6 even without cheating.

Similarly, why would subjects cheat on “3” but not on “4”, given that both would yield 1 point? There is also some weak evidence for cheating on “2” (one-tailed binomial test, $p = 0.1$), but not on “5”. Although more data are needed to confirm this pattern, it hints at the possibility that subjects prefer to cheat by choosing “D” in the face of “2” and “3” in Transparent. The proportion of “D” chosen by the two cheaters turns out to be significantly higher than that by non-cheaters in Transparent (MWW test, $p = 0.05$), while no difference of this sort is found in Opaque. This is a bit surprising because one would expect that it is more difficult to forego the gains when seeing them, i.e., that one is more tempted to untruthfully report “U” when one sees “4” than reporting “D” when one sees “3”. One possible explanation is that reversing the side for obtaining the higher earning can be relatively more indirect and thus subjects are less directly confronted with the act of cheating, especially when decisions are made quickly.

4.2 Opacity and Reaction Time

The results in this section are based on reaction time data, assuming that more reaction time indirectly indicates more cognitive efforts exerted (e.g., Rubinstein, 2007). The main question concerns whether different cheating contexts result in more or less cognitive efforts depending on subjects’ types¹⁷. For instance, if it holds that the opacity of the cheating context allows for more self-deceptive excuses, which in turn should require less cognitive effort while cheating, then the time spent in each round in Opaque should be shorter than in Transparent for the dishonest. As for the honest, the cognitive efforts exerted should not differ according to the recent findings of Greene and Paxton (2009) which conclude that honesty results more from the absence of temptation rather than the active resistance of temptation.

Table 6 displays the main results: in general, as shown in column A, subjects on average spent about 11 seconds per round in Transparent, but 3 seconds more in Opaque. When taking into account the interaction effects, as shown in column B, the dishonest in Opaque turns out to spend significantly less time than either the honest in Opaque by almost 4 seconds, or the dishonest in

¹⁷ Since the exact classifications according to the two measurements do not fully overlap, subjects are defined as “honest” or “dishonest” only if they are classified consistently according to both measurements. As a result, 25 subjects are defined as “honest”, 7 as “dishonest” and the other 7 subjects are classified as “ambiguous” because of conflicting classifications.

Transparent by 2 seconds. Though the result is perfectly in line with the supposition that the dishonest subjects consume less cognitive efforts for cheating per round, it is however not conclusive given that we only have one common identifiable cheater in Transparent according to both foresight and earnings. This subject could well be an outlier. Nevertheless, it would be an interesting follow-up to examine if the availability of self-deceptive excuses in the cheating context leads to less cognitive efforts or less time spent on cheating.

Table 6: OLS Regression Results on Reaction Time

Reaction Time (Seconds per round)	A	B	C
Opaque	3.27** (1.31)	5.89*** (2.03)	5.11** (2.09)
Ambiguous		-2.25* (1.30)	-3.03** (1.39)
Dishonest		4.33*** (0.71)	3.55*** (0.86)
Ambiguous * Opaque		-4.13* (2.39)	-3.36 (2.44)
Dishonest * Opaque		-8.19*** (2.23)	-7.42*** (2.28)
Using Pattern			-2.33* (1.25)
Constant	11.20*** (0.67)	11.20*** (0.71)	11.98*** (0.862)
R-squared	0.02	0.05	0.05

Note: *, **, *** indicates significance at the 90%, 95%, and 99% level, respectively; standard errors are reported in parentheses and clustered on each subject.

Reaction time per round is calculated as the difference between the starting time of every two rounds which is automatically recorded online.

There are 7 subjects who are classified as “Ambiguous”, 7 “Dishonest” and 25 “Honest”.

Number of Observations ¹⁸: 741

As for the honest subjects, I found a reverse pattern. While the honest subjects in Transparent spent only 11 seconds per round, the honest subjects in Opaque spent 6 seconds more. This seems to suggest that subjects struggle more in Opaque to be honest. Moreover, I look further to see whether the difference is merely driven by the pattern users in Transparent who decide only once

¹⁸ There is no round time data of the last round since subjects were asked to visit the survey webpage without necessarily closing down the experiment webpage.

for all 20 sides. As is shown in the last column, the subjects using patterns in Transparent indeed spent on average 2.3 seconds less than the rest. But the treatment difference is still present, controlling for pattern-using ($p < 0.05$), especially given that 5.1 seconds is a considerable difference.

Yet another possible explanation is as follows: as subjects in Opaque were not asked to write down the side immediately after the choice in mind, the recalling of the side can require honest subjects additional cognitive efforts. But, since a round lasts less than 20 seconds and the most time-consuming step is writing down the side, it is unlikely that it takes 5 seconds more to recall the side in Opaque than in Transparent.

The preliminary result that honest subjects spend more time in Opaque potentially implies that the cognitive efforts for honesty can differ corresponding to the level of opacity in different cheating environments: when it becomes sufficiently opaque, it can require more active resistance and thus more efforts to be honest, whereas when the cheating context is transparent, it encourages subjects to pre-commit to honesty and saves subjects from the trap of temptation. Since Greene and Paxton (2009)'s results were based on only one cheating context, more studies varying the level of opacity should be carried out before settling down the ongoing debate of whether honesty results from the active resistance or the absence of temptation.

Since subjects spent much more time in Round 1 than the rest, I also ran the regressions without Round 1. The results are robust for the comparison regarding the honest subjects as well as the within treatment comparisons, as also demonstrated by the tabulations in Table 7.

Table 7: Round Time

Round time (seconds)	Round 1		Round 2-19		All	
	Honest	Dishonest	Honest	Dishonest	Honest	Dishonest
Opaque	45.5	45.7	15.5	11.4	17.1	13.2
Transparent	36.1	27.8	9.8	14.9	11.2	15.5

I also compare the time spent on reporting “U” or “D”. The time used reporting either side is the same in Opaque. However, when the round foresight equals to one, non-pattern-users in Transparent spent, on average, 4 seconds more when reporting “D” (t-test, $p = 0.03$). This is

probably not because reporting “D” generally takes more time, since no difference of the sort is found when foresight is zero. Rather, this is likely to be another indication that cheating in Transparent requires more cognitive efforts given that subjects cheat more by reporting “U” (see *Result 4*).

In addition, subjects spent 1.6 seconds more on average in the honest rounds ($f = 0$) than the dishonest or unlucky rounds ($f = 1$) in Opaque (t-test, $p = 0.09$), while no difference is found in Transparent.

All in all, the reaction time results seem to suggest that opaqueness can potentially trigger different psychological processes in resisting or giving in to the cheating temptation. More specifically, the overall results suggest that it requires more effort to be honest in Opaque than in Transparent, and it is more effortful to be honest than to cheat in Opaque.

5. Discussion

Although the theory of self-image maintenance seems to have provided key insights into cheating behavior of honest people, more understandings are pending regarding how this theory can account for the intricate cheating patterns shown in the experimental literature. I argue that taking an evolutionary perspective can help us understand what the crucial factors of self-image are and ultimately shed new light on how people cheat. In particular, the experiment in this paper highlights the important role that intentions play in self-image maintenance and cheating.

Trivers (1971) argues that good intentions are crucial for choosing a good partner for a long-term relationship of reciprocal helping. To cooperate with the “right” counterparts, humans are engaged in the game of signaling (detecting) good (bad) dispositions or intentions. For detecting cheaters, it is the intent or motive behind the act that can more reliably signal the altruistic disposition. The intent behind an act is more important than the act itself for two main reasons. First, one can trust someone with a genuinely good intent to behave cooperatively in the long run even without any monitoring. Second, one does not want to falsely exclude a cooperator who accidentally caused an error in outcome. There are legitimate excuses for accidental bad outcomes that can go beyond the will of an agent, but not for accidental bad intents. In turn, intent detection

has also triggered the adaptive response of signaling good intents or disguising bad intents so that people must evolve better tricks to better detect fake intentions. As pointed out by von Hippel and Trivers (2011), even self-deception can be understood as another adaption to better hide deceptive intents in response to the evolved detection of sham intents. Likewise, our caring to be moral or maintaining a moral self-image can potentially also be a result of the arms-race between covering and uncovering cheating (Jiang & Lindemans, 2012).

If moral self-image maintenance is essentially part of the morality game, the self-image that we strive to maintain is unlikely to be one of not making any error or mistake in outcomes, but one of having good intents. In this way, “cheating only a little bit” can also be explained as a strategy to mimic cheating outcomes that result from errors instead of the intent to cheat. Cheating a lot is thus not a good strategy because it cannot be easily accounted by ignorance or errors since big accidental errors are usually rarer in daily life.

This study argues that people cheat more when they can hide their bad intentions with self-deceptive excuses. For those who care to be moral, they will feel less guilty if they manage to deceive themselves that their intentions are good. I argue that subjects cheat less in Transparent because it is more difficult to deceive oneself there that the intent to cheat is absent.

There might be a few alternative explanations. First, since subjects have to break an additional rule in Transparent, it is possible that subjects cheat less because they are averse to rule-breaking *per se* and, as a result, face an additional psychological barrier to cheat. Some subjects, however, do not seem to mind breaking certain rules of the game in general: one of the subjects openly reported her rule-breaking in the post-experimental survey, “In the first several rounds, I decided randomly. After that, I filled out all the rest of the form.” We probably would not think that this person is morally bad. We might even perceive him as morally good since he or she probably did not intend to cheat. On the other hand, it is hard to imagine that subjects would openly report that they switched the order of the steps since the rule breaking there can be linked to cheating.

Nevertheless, this seems to be again related to intentions: subjects fear any exposure of breaking the rule of the order of the game because it would reveal the cheating intent, even though again nothing could be said about whether they had truthfully reported the side. This points to the possibility that breaking a rule is not judged as morally bad *per se*. It also raises an interesting question of whether we view lying in our daily life in a similar way. As long as we believe or

deceive ourselves to believe that our intentions behind those lies are not bad, we can simply lie without feeling guilt or shame.

Relatedly, it is still possible that subjects cheat less in Transparent out of concern for a good moral reputation instead of moral self-image, if they are simply afraid of being suspected (even though not caught) of having cheated despite the double blind payment procedure. The probability of being suspected would be higher in Transparent if one is also caught switching the steps. However, it is not clear how these two mechanisms can be clearly separated as they are often interwoven. Nevertheless, it would be an intriguing follow-up exercise to examine the different cheating patterns in deceiving others and self.

After all, this study posits that the opaqueness of the cheating context affect how much people cheat and how they cheat. The conjecture is closely related to existing theories. For instance, one of the key factors that Mazar et al. (2008) suggests is “categorical malleability” which refers to the ease of categorizing actions in a self-serving manner. It is plausible that the degree of malleability of a cheating context is determined by the ease of intent inference. Fischbacher and Heusi (2008) postulate that subjects avoid cheating for the maximal possible benefit in order to show that they are not greedy. The desire to signal a non-greedy image again seems to be related to the omnipresent need of our signaling a good intention, even when we cheat.

6. Conclusion

By exploiting two novel variants of the mind game, I show in this paper how a subtlety in the rules of the game significantly affects cheating behavior. The overall results strongly suggest that the more inferable the intent to cheat, the less people cheat. Whether people have cheating in mind can be more important than the outcome of the cheating act. Moreover, intricate cheating patterns also seem to systematically differ in these two cheating contexts, potentially triggered by the ease of inferring the intent the cheat.

Though the role of intent detection in deterring cheating is not obvious, it is intuitive that subtle cheating requires subtle deterrent mechanisms. It cannot be overemphasized the importance of taking into account intent detection in designing rules that save individuals from the “trap” of

cheating. More specifically, policy makers should not only increase transparency in overseeing the potential rule-breaking acts, but also try to make sure that the rules minimize the possibility of unintentional cheating. Rules should be designed so that individuals are confronted with their own bad intents if they cheat. More specifically, rules should be designed so that cheating requires additional explicit steps that can clearly reveal the intent to cheat. In this way, bad intents can no longer be covered up fuzzy contexts and self-deceptive excuses, but instead get exposed to the light of humans' own conscience. Moreover, opaque rules that leave room for the excuses of "unintentional cheating" also shift the burden of proof away from the rule-breakers, in accordance with the legal principle of "*in dubio pro reo*" (*when in doubt, [judge] in favor of the accused*). Above all, the cost of using external enforcement mechanisms, such as increasing the probability of cheating detection, is notoriously high, let alone that these mechanisms are not always applicable to those who care about their moral self-image. In contrast, intent detection is a costless deterrent mechanism that relies on moral self-regulation without over-burdening the legal system.

New questions that emerged from the results remain to be addressed in future work. It would also be interesting to further examine how cheating responds to the level of vagueness of the context since vagueness is crucial for self-deception (Sloman et al., 2010). In this experiment, questions can be inserted to make it harder for subjects to rely on the excuse of errors or ignorance, such as, "In case you do not remember clearly which side you have chosen, please choose again." Also as a next step, neuroscientific studies can be carried out to examine the intent related neural activities in these two treatments. Above all, the mind game a useful tool especially in light of the complexity and subtlety of the cheating patterns of the honest.

References

- Abeler, J., Becker, A. & Falk, A. (2012). Truth-telling – A Representative Assessment. *Working Paper*
- Ariely, D (2008). *Predictably Irrational*, New York: Harper Collins.
- Bandura, A. (1990). Selective Activation and Disengagement of Moral Control. *Journal of Social Issues*, 46(1), 27-46.
- Bandura, A., Barbaranelli, C., Caprara, G. V., & Pastorelli, C. (1996). Mechanisms of Moral Disengagement in the Exercise of Moral Agency. *Journal of Personality and Social Psychology*, 71(2), 364-374.
- Becker, G. S. (1968). Crime and Punishment: An Economic Approach. *Journal of Political Economy*, 76(2), 169.
- Benabou, R., & Tirole, J. (2002). Self-Confidence and Personal Motivation. *The Quarterly Journal of Economics*, 117(3), 871-915.
- Chance, Z., Norton, M. I., Gino, F., & Ariely, D. (2011). Temporal View of the Costs and Benefits of Self-deception. *PNAS*, 108(3), 15655-9
- Charness, Gary and Dufwenberg, Martin. (2006). Promises and Partnership. *Econometrica*, 74(6): 1579-601
- Croson, Rachel T.A. (2005). "Deception in Economic Experiments," In *Deception in Markets: An Economic Analysis*. Ed. Caroline Gerschlager, 113-130. New York: Palgrave Macmillan
- DePaulo, B. M., Kashy, D. a, Kirkendol, S. E., Wyer, M. M., & Epstein, J. a. (1996). Lying in Everyday Life. *Journal of personality and social psychology*, 70(5), 979-95.
- Fischbacher, U., & Heusi, F. (2008). Learning and Peer Effects Lies in Disguise: An Experimental Study on Cheating, *TWI Research Paper Series (40)*
- Gneezy, Uri. 2005. "Deception: The Role of Consequences." *The American Economic Review*, 95: 384-94.
- Greene, J. D., & Paxton, J. M. (2009). Patterns of Neural Activity Associated with Honest and Dishonest Moral Decisions. *PNAS*, 106(30), 12506-11.
- Houser, D., Vetter, S. & Winter, J. K. (2011). Fairness and Cheating. *GMU Working Paper in Economics*, 11-02

- Jiang, T. (2012). Whom Do you Trust in a Bribery Game? *Working Paper*
- Jiang, T. & Lindemans, J. (2012). Weird Behavior in Cheating Games: The Coevolution of Deception and Detection. *Working Paper*
- Lundquist, T., Ellingsen, T., Gribbe, E., & Johannesson, M. (2009). The Aversion to Lying. *Journal of Economic Behavior & Organization*, 70(1-2), 81-92.
- Mazar, N., Amir, O., & Ariely, D. (2008). The Dishonesty of Honest People: A Theory of Self-Concept Maintenance. *Journal of Marketing Research*, 45(6), 633-644.
- Mead, N. L., Baumeister, R. F., Gino, F., Schweitzer, M. E., & Ariely, D. (2009). Too Tired to Tell the Truth: Self-Control Resource Depletion and Dishonesty. *Journal of experimental social psychology*, 45(3), 594-597
- Rubinstein, A. (2007). Instinctive And Cognitive Reasoning: A Study Of Response Times. *The Economic Journal*, 117, 1243-1259.
- Schweitzer, M., Hershey, J., & Bradlow, E. (2006). Promises and Lies: Restoring Violated Trust. *Organizational Behavior and Human Decision Processes*, 101(1), 1-19.
- Shu, L. L., Gino, F., & Bazerman, M. H. (2011). Dishonest Deed, Clear Conscience: When Cheating Leads to Moral Disengagement and Motivated Forgetting. *Personality & social psychology bulletin*, 37(3), 330-49.
- Sloman, S.A., Fernbach, P. M., & Hagmayer, Y. (2010). Self-deception Requires Vagueness. *Cognition*, 115(2), 268-81.
- Shalvi, S., Handgraaf, M. J. J. and De Dreu, C. K. (2011), Ethical Manoeuvring: Why People Avoid Both Major and Minor Lies. *British Journal of Management*, 22: S16–S27.
- Trivers, R.L. (1971). Evolution of Reciprocal Altruism. *The Quarterly Review of Biology*, 46(1), 35-57.
- Von Hippel, W., & Trivers, R. (2011). The Evolution and Psychology of Self-deception. *The Behavioral and brain sciences*, 34(1), 1-16; discussion 16-56.

Appendix I: Instructions (Treatment Opaque)

Introduction:

It is known that human beings seem to expect dependencies between successive events in spite of the fact that they know that the events occur independently of each other. Individuals' successive responses tend to be mutually dependent towards outcomes that are randomly generated. This experiment studies individuals' successive responses towards outcomes that are randomly generated in an incentivized lottery setting.

Payment:

At the beginning of the experiment, you will be asked to come up with a personal password for the sake of the payment. As of 20th September, you can pick up your payment by presenting your password to the secretaries of Department of Economics in K412. At the end of the experiment, you will be also asked to fill in a short online questionnaire in which you have to fill in the same password as the one for the experiment in order to guarantee an additional reward for the questionnaire.

Die-throwing game: Instructions

You are about to play a die-throwing game. In this game, you can throw a virtual online die 20 times for earning points. All resulting points will be exchanged to Euros. Every point you earn is equivalent to 5 Eurocents.

The die has six sides and each side has a different number of dots. The pairs of numbers add up to 7 on the opposite sides: 1 vs. 6, 2 vs. 5 and 3 vs. 4, and vice versa. We call the visible side facing up of any dice you throw the up side "U" and the opposite invisible side which faces down the down side "D". For instance, if the up side is 1, the down side of the die is 6 (See the table below for all the sides.)

U	1	2	3	4	5	6
D	6	5	4	3	2	1

In each round, the number of points that you score, depends on the throw of the die as well as on the side that you have chosen in that round. Each round consists of one throw. Before throwing, you have to choose the relevant side for that round. Note that the die outcomes are random and the outcome you see on the screen corresponds to the upside. We will explain you the detailed procedure of each round in steps:

Step 1: before every throw, make a choice of side in your mind between “U” (Up) and “D” (Down).

Step 2: throw the virtual die by clicking on the “Throw Die” button on the screen.

Step 3: fill in the side (“U” or “D”) you have chosen in your mind in Step 1 in the outcome form.

On the basis of the outcomes in step 2 stored on our web server and the corresponding side filled in the outcome form, your final earning points will be calculated and converted to Euros as your payment. For instance, if you have chosen “D” in your mind and the die outcome turns up to be “4”, you earn 3 points for that throw, whereas if you have chosen “U” in your mind, you earn 4 points.

Test questions:

1. If you chose side “D” in your mind and the die outcome turns up to be “5”, how many points do you earn?_____
2. If you chose side “U” in your mind and the die outcome turns up to be “6”, how many points do you earn?_____

Appendix II:

Outcome Form (Treatment Opaque)

Please fill in your personal password for picking up your payment: _____.

Please repeat the following three steps for each of the 20 Throws:

Step 1: before each throw, make a choice of side in your mind between “U” (Up) and “D” (Down).

Step 2: throw the virtual die by clicking on the “throw die” button on the screen.

Step 3: fill in the side (“U” or “D”) you have chosen in your mind in Step 1 in the outcome table.

	T ₁	T ₂	T ₃	T ₄	T ₅	T ₆	T ₇	T ₈	T ₉	T ₁₀	T ₁₁	T ₁₂	T ₁₃	T ₁₄	T ₁₅	T ₁₆	T ₁₇	T ₁₈	T ₁₉	T ₂₀
The Side “U” or “D”																				

When you are finished with all the steps, please put the outcome form in the envelope and open the website under Internet Explorer’s favorite menu - “a short questionnaire”. When you are finished with the questionnaire, please turn off all IE windows and inform the experimenter.

Outcome Form (Treatment Transparent)

Please fill in your personal password for picking up your payment: _____.

Please repeat the following three steps for each of the 20 Throws:

Step 1: before each throw, make a choice of side in your mind between “U” (Up) and “D” (Down).

Step 2: fill in the side (“U” or “D”) you have chosen in your mind in Step 1 in the outcome table.

Step 3: throw the virtual die by clicking on the “throw die” button on the screen.

	T ₁	T ₂	T ₃	T ₄	T ₅	T ₆	T ₇	T ₈	T ₉	T ₁₀	T ₁₁	T ₁₂	T ₁₃	T ₁₄	T ₁₅	T ₁₆	T ₁₇	T ₁₈	T ₁₉	T ₂₀
The Side “U” or “D”																				

When you are finished with all the steps, please put the outcome form in the envelope and open the website under Internet Explorer’s favorite menu - “a short questionnaire”. When you are finished with the questionnaire, please turn off all IE windows and inform the experimenter.

Appendix III: Screen Shots

Treatment Opaque

Screen Start:

If you are ready, please start the game by filling in your personal password: _____

Screen 1:

Round X

If you have chosen the side (“U” or “D”) in your mind, please click on the button below:

Throw the Die

Screen 2:

The die outcome for this round is:



Next Round

Screen Final:

You have finished all the 20 rounds. You can proceed by clicking on "a short questionnaire" under the IE “Favorites” menu.

Treatment Transparent

Screen Start:

If you are ready, please start the game by filling in your personal password: _____

Screen 1:

Round X

If you have chosen the side (“U” or “D”) in your mind, please fill in the side on the outcome form and click on the button below:

Throw the Die

Screen 2:

The die outcome for this round is:



Next Round

Screen Final:

You have finished all the 20 rounds. You can proceed by clicking on "a short questionnaire" under the IE “Favorites” menu.