1. INTRODUCTION

People often cheat (e.g. Herman, 2005; Mazar et al., 2008). Most often, they do so in order to amass wealth, attain power, avoid adverse consequences, or otherwise benefit themselves (e.g. Gneezy, 2005; see a review by Alm and McKee, 1998 on tax evasion). There is ample evidence that cheating among university students is rampant (e.g. McCabe and Trevino, 1993; 1997). Cheating is also common in the business world (e.g., Buffet, 2003; Jensen, 2001, 2003). According to a recent publication by the Association of Certified Fraud Examiners (ACFE, 2014) US businesses lose approximately 5% of their annual revenues to various forms of consumer fraud, an amount equal to at least $1 trillion in losses across the economy. Researchers have analyzed the personal characteristics that are associated with cheating as well as the situations that are most likely to tempt people to cheat (e.g., Bandura, 1999; Becker, 1968; Cadsby et al., 2010; Chugh et al., 2005; Ehrlich, 1973; Erat and Gneezy, 2012; Gneezy, 2005; Nagin et al., 2002; Jones, 1991; Reynolds, 2006; Song and Zhong, 2015).

Cheating is generally regarded as unethical and is associated with moral failure. However, not all cheating is purely selfish. Sometimes people cheat out of their concern for others, particularly if those people are family, friends, classmates, work colleagues or team members (e.g., Conrads et al., 2011; Diekmann, et al., 1997; Gino and Pierce, 2010; Schweitzer and Hsee, 2002; Wiltermuth, 2011). A university student lends an essay to a classmate to copy. A soldier or policeman hides the full truth in order to protect a colleague who has used excessive force on an enemy soldier or suspected criminal. A skating judge ensures the team from his/her own country wins the competition. Generosity and altruism are usually considered virtues. It seems paradoxical that such virtues may sometimes be associated with cheating, which is normally considered a vice.

Such cheating for others is not harmless. By cheating on behalf of a member of one's own social group, one is often hurting one or more people outside that group, causing a negative externality to a third party (e.g., Becker, 1957; Bernhard et al., 2006a, 2006b). By giving the job to the niece of a current employee, the employer may be passing over more qualified applicants, which is costly both for the organization and for applicants without the necessary connections. In this paper, we present a controlled laboratory experiment, which allows us to investigate whether, and to what extent, people will cheat on behalf of a member of their own social group at the expense of a non-member. Specifically, we examine whether moral concern constrains people from cheating in favor of a member of their own group and whether the moral burden of cheating is as strong a deterrent for such cheating as it is for selfish cheating.

To investigate the impact of social/group identity on cheating, we conduct a money-allocation task in the laboratory, varying two factors: (1) whether the allocator is permitted to choose his/her preferred allocation freely from the available allocations, or is required to report the results of a private die-roll that determines the allocation, coupled with an opportunity to cheat (Fischbacher and Föllmi-Heusi, 2013); (2) whether the allocator allocates money either between an in-group person and an out-group person, or between him/herself and an out-group person. Varying the first factor allows us to compare revealed allocation preferences with those that emerge under a process that should be free of bias unless systematic cheating occurs. Varying the second factor permits us to compare revealed allocation preferences, cheating behavior, and the extent to which the moral burden of cheating inhibits the pursuit of revealed allocation preferences under the in-group/out-group versus the self/other condition.

We find moderate but significant in-group bias when the allocator can freely choose the allocation between an in-group and an out-group member. If people follow the prescribed rules, the Die-Roll condition should eliminate such a biased allocation. While the distribution of allocations changes significantly under the Die-Roll condition, the moral burden of cheating for others has no significant impact on the average amount allocated to an in-group member, which is statistically indistinguishable from the average amount allocated under the Free-Will condition.

This stands in marked contrast to the outcome when the allocation is between oneself and an out-group person. When one can freely choose this allocation, the average amount allocated to oneself is significantly greater than the average amount allocated to an in-group member under the in-group/out-group condition. However, the average amount allocated to oneself in the Die-

Roll condition, while still biased toward oneself, is significantly lower than in the Free-Will condition. This suggests that the moral burden of selfish cheating, while failing to produce a completely honest outcome, does place some constraint on average self-interested allocations. Despite the exhibited restraint, allocators still lie sufficiently to take more for themselves on average than they allocate to an in-group member in the in-group/out-group Free-Will or Die-Roll treatments.

The remainder of this paper is structured as follows. Section 2 discusses the related literature. Section 3 presents the details of our experimental design. Section 4 outlines behavioral predictions. Results are in section 5, and we conclude in section 6.

2. RELATED LITERATURE

Social identity theory, developed by Tajfel and Turner (1979), is a social psychological theory of group processes, intergroup relations and the basis of social self. It has been developed and applied to understand social phenomena such as racism, prejudice and intergroup discrimination. The basic idea is that people tend to perceive themselves and others through various social categories, such as organizational membership, religious affiliation, gender, and age cohort, and that these perceptions can motivate in-group favoritism, namely the preferential treatment of in-group over out-group members (Tajfel and Turner, 1985).

Early experimental work in social psychology developed the methodology of minimal group paradigm (MGP) to identify the emergence of in-group favoritism (Tajfel et al, 1971). Specifically, subjects are randomly divided into two groups ostensibly based on a trivial criterion such as a preference for paintings by Klee or Kandinski. Tajfel and Turner (1985) reviewed the considerable experimental research using the MGP up to the mid 1980s, and concluded that social categorization alone was sufficient to produce in-group favoritism.

Social identity theory was formally incorporated into economic theory by Akerlof and Kranton (2000), who supplemented the standard economic framework by recognizing that individual choices and behaviors are influenced by social identity. In their framework, identity enters the utility function through its prescription of behavioral norms, the departure from which causes disutility commensurate with the extent of the departure. This framework has been applied to yield new insights into a variety of socioeconomic phenomena such as gender discrimination (Akerlof and Kranton, 2000), the household division of labor (Akerlof and Kranton, 2000), the economics of education (Akerlof and Kranton 2002), economic development (Basu, 2006), trust (Buchan et al., 2002; Buchan et al., 2006), public goods provision (Solow and Kirkwood, 2002; Croson, et al., 2008; Eckel and Grossman, 2005), preferences over redistributive tax regimes (Klor and Shayo, 2010) and cooperation (Ruffle and Sosis, 2006; McLeish and Oxoby, 2011).

While some experiments by economists have employed a near or enhanced minimal group paradigm[1] to induce group identity within the lab while eschewing deception (e.g., Buchan et al., 2002; Eckel and Grossman, 2005; Buchan et al., 2006; Chen and Li, 2009; Chen and Chen, 2011; Harris et al., 2015), others have used natural group identities (e.g., Croson et al., 2008; Ruffle and Sosis, 2006; Klor and Shayo, 2010; McLeish and Oxoby, 2011; Ockenfels and Werner, 2014). At least one study used both (Solow and Kirkwood, 2002).

Several recent studies are particularly notable. First, Yan Chen and colleagues (Chen and Li, 2009; Chen and Chen, 2011) conducted a series of lab experiments demonstrating that under experimentally induced identities mere categorization into groups can produce group-contingent social preferences and affect equilibrium selection. Moreover, they showed that subjects, when interacting with an in-group member, exhibited significantly more altruism, reciprocity, forgiveness and overall efficiency maximization than when interacting with an out-group member. Second, Ockenfels and Werner (2014) used natural groups, i.e. the student subject's university, to demonstrate that in-group favoritism is partly belief-dependent. In particular, subjects in a dictator game transferred significantly more to in-group recipients whom they knew

[1] See Chen and Chen (2011) for a detailed discussion of these terms. Methods used by economists to induce identity in the lab include transparent random assignment to artificial groups without any pretense that the assignment is based on even a trivial criterion (e.g. Chen and Li, 2009; Chen and Chen, 2011; Harris et al., 2015), non-random assignment actually (i.e. non-deceptively) based on a trivial criterion such as a preference for one artist over another (e.g. Chen and Li, 2009), and random assignment enhanced by a group-building exercise such as performing a task together (e.g. Solow and Kirkwood, 2002; Buchan et al., 2002; Eckel and Grossman, 2005; Buchan et al., 2006; Chen and Chen, 2011).

were aware of the shared group identity than to in-group recipients whom they knew were unaware of the shared identity. Third, Harris et al. (2015) examined whether in-group favoritism is regarded as a social norm or a violation of social norms using transparent random assignment to minimal groups in the context of a dictator game followed by a punishment opportunity. The authors found that perception and judgment of in-group favoritism depends on the perceiver's social/group identity. Specifically, inferring from punishment behavior, in-group favoritism is acceptable to in-group members, unacceptable to out-group members, and of little consequence to third parties, implying that in-group favoritism may be a social norm in its own right rather than a violation of norms among in-groups.

In this paper, we investigate the following question: when in-group favoritism is present, will people cross a moral boundary by lying and cheating to help in-group members? Recent papers in behavioral and experimental economics have examined cheating both theoretically and empirically. For example, Gneezy and colleagues (Erat and Gneezy, 2012; Gneezy, 2005) examined how the consequences of lying affect its likelihood, and find that the higher the self-interested benefit from lying and/or the lower the cost for the deceived party, the greater the likelihood that a person will choose to lie. Charness and Dufwenberg (2006) noted that guilt-aversion is an important inhibitor to cheating behavior. Schweitzer et al. (2004) showed that people lie more frequently when only slight exaggeration is required to meet a pre-announced target, while Cadsby et al. (2010) demonstrated that target-based compensation systems are associated with a greater likelihood of misrepresentation than piece rates or tournaments. This is because the trade-off between the size of the required lie and the resultant financial reward is usually lower in the former case. Small lies are apparently less psychologically costly then big ones. Other researchers have shown that people are more willing to lie when not only the liar, but others as well, benefit from the dishonest behavior (Conrads et al., 2013; Diekmann, et al., 1997; Gino and Pierce, 2010; Schweitzer and Hsee, 2002; Wiltermuth, 2011).

Fischbacher and Föllmi-Heusi (2013) developed a simple and elegant private die-rolling paradigm to study cheating behavior. Under this "die-under-cup" paradigm, which allows participants to report the outcome of a die roll only they can see and gain money according to their reports, the authors find that people systematically over-report the outcome of the private die-roll when their payoff is determined by it. However, most people are only "partial" cheaters in the sense that they do not report the die outcome that would maximize their earnings. This is in line with the idea of "self-concept maintenance" suggested by earlier work in this area (e.g. Shalvi et al., 2011, 2012; Shalvi and Leiser, 2013; Gino and Ariely, 2012; Mazar, et al., 2008; Houser et al., 2012).

After completing our study, we became aware of two recent studies that like ours examine cheating on behalf of one's in-group, but in each case with a rather different focus. The first study, Jiang (2014), focuses on inter-country differences in cheating behavior. Using samples from China, Japan, Italy and the Netherlands and random assignment to artificial groups coupled with a team-building exercise, Jiang (2014) finds that cheating for self and for another varies with a country's corruption ranking rather than with the extent to which it is deemed to be a collective versus an individualistic society. Opportunities to ignore a pre-specified allocation rule in favor of cheating are always at the expense of the experimenter rather than another subject in the experiment. Among Chinese university students, there is evidence of such cheating for oneself, for an in-group member, and for an out-group member. However, the authors do not report on whether the amount of cheating differs significantly depending on the beneficiary of the cheating.

The other study, Hruschka et al. (2014), utilizes samples from eight societies in six countries to investigate whether a greater threat of infectious disease and/or the relative absence of institutions that allow basic needs to be met through impartial interactions may be associated with a willingness to ignore impartial allocation rules to exhibit in-group favoritism. They employ two treatments. In one, an allocator is instructed to follow a specified rule that allocates a resource randomly between the allocator and an out-group person. In the other, the allocator is instructed to follow the same rule to divide a resource between another in-group person and an out-group person. However, in both treatments the allocator can ignore the rule with no chance of detection. The results suggest that it is the lack of effective institutions rather than the greater threat of infectious disease that is associated with ignoring the allocation rule to favor oneself or one's in-group member. One of Hruschka et al.'s (2014) samples uses Chinese university students and natural groups as we do in our study. In contrast to Jiang (2014), they find no evidence of the Chinese students ignoring the allocation rule to benefit either themselves or their in-group members. We return to both these papers to compare their methodological approaches and results with ours in our concluding discussion.

3. EXPERIMENTAL DESIGN

3.1 Induced or Natural Groups

An important design choice was whether to attempt the creation of in-groups within the lab or to use natural in-groups. As indicated above, there are many studies of both types in the experimental economics literature. Each approach has advantages. The successful induction of group identity through the minimal group paradigm has a long tradition in psychology. It gives the experimenter control that may be difficult or lacking when individuals have a multitude of natural identities (Chen and Chen, 2011). It allows the experimenter to strengthen or weaken the group-building manipulation (Chen and Chen, 2011). Most importantly, it is easy to combine with random assignment to groups, which is necessary to infer unambiguously a causal relationship from an observed correlation.

Naturally occurring groups, however, may be more directly relevant to day-to-day behavior outside the lab. Along these lines, Klor and Shayo (2010, p. 272) argue that using natural rather than artificial groups ensures "that groups have some meaning outside of the laboratory." While natural groups occur in every culture, induced groups are not necessarily created with equal ease and effectiveness in every culture. Triandis et al. (1988) argues that people in more individualist cultures may be better at meeting new people and forming new in-groups than people in more collectivist cultures. This suggests that it may be more difficult to form new in-groups in the lab in such a cultural context. Brewer and Yuki (2007, p. 317) argue that in-group favoritism based on the arbitrary category distinctions that characterize groups created by the minimal group paradigm may be “more pronounced in Western than in Asian cultures.”

Our experiment is set in China, an Asian society often deemed to be collectivist in nature. We know of four experimental economics studies that use induced groups in the Chinese context: Buchan et al. (2002), Buchan et al. (2006), Jiang (2014), and Eriksson et al. (2015). Buchan et al. (2002) studied an indirect-exchange variant of the trust game in the United States, Korea, Japan, and China. Citing Triandis et al. (1988), they hypothesize that their minimal-group induction technique will be effective at creating more cooperation among individually-oriented individuals, but not among collectively-oriented individuals. This hypothesis is supported by their data, suggesting caution before implementing such a technique to create in-group favoritism in a country such as China with a collectively-oriented culture. Buchan et al. (2006) implement a standard trust game in the same four countries. They find that in the US, participants both send and return more to induced in-groups than to induced out-groups, while in China more is both sent and returned to induced out-groups than to induced in-groups. A cultural-orientation questionnaire reveals that Chinese participants are more collectively oriented than American participants, and a mediation analysis demonstrates that cultural orientation fully mediates the observed differences in behavior between Chinese and Americans. For our purposes, the important conclusion is that the methods of inducing group identity that have worked well to produce in-group favoritism in the US and other Western countries may fail to do so in the different cultural environment of China.

Jiang (2014), described above, is unable to find support for her hypothesis that subjects oriented toward strong in-group favoritism should cheat more. In fact, she finds that they cheat less. She suggests (Jiang, 2014, p. 26) that this is because the artificial groups used in her experiment are most effective at creating in-group favoritism for those subjects who “are more universalistic in real life and open to make new ties with strangers.” She thus urges caution in “interpreting or generalizing” her “related in-group bias results” and suggests the gathering of real in-group data to shed light on this issue. It is unclear whether Jiang's results for China exhibit significant in-group bias.

Eriksson et al. (2015) investigate whether people are willing to sacrifice some of their own resources to save the face of another. Their experiment is run in China, and groups are induced by an experimental manipulation. They find that their subjects are indeed willing to pay to save someone else from embarrassment, regardless of whether that person is an in-group or out-group member. The authors observe that this may be because saving face is a general social norm that is unrelated to social distance. However, it may also reflect the possibility that the minimal-group induction process was less effective in China than it might be in a more individualistically oriented country such as the US.

Our study requires an environment that exhibits in-group favoritism. Without the presence of in-group favoritism, it is impossible to investigate whether people might cheat on behalf of a member of the in-group they favor. Given the lack of evidence that in-group favoritism can be successfully induced in China, we turn to naturally occurring groups as our experimental manipulation of social identity. Specifically, following Ockenfels and Werner (2014), our subjects were students enrolled at two different universities in nearby cities: Zhejiang University (ZJU) and Shanghai University of Finance and Economics (SUFE). Both of these universities are among the top universities in China with highly competitive admission standards. ZJU is a very large comprehensive university with particular strength in science located in Hangzhou, the capital of Zhejiang province 163 kilometers (101 miles) from Shanghai. SUFE is a smaller university specializing in economics, finance and accounting. Based on the most recent ranking of Chinese universities in 2015, ZJU is #8 and SUFE is #10 among all universities in China in the list of "most favored universities by potential students" (www.cuaa.net), and in terms of "graduate salary ranking", ZJU is #12 and SUFE is #3 among all universities in China (www.biaozhun007.com).

A benefit of choosing nearby universities with comparable academic standards and similar demographics is that we remove possible confounds such as differing socio-economic status or academic ability to focus on social identities arising purely from the group to which one belongs represented in our experiment by the university that one attends. While this differs from the minimal group paradigm because identity is not created in the lab, it may be thought of as a minimal naturally-occurring group manipulation similar to that used successfully by Song et al. (2012) in their investigation of trust, reciprocity and in-groups in China. We discuss the generalizability of our results based on natural groups in the concluding discussion.

3.2 Experimental Procedures

To observe and measure in-group favoritism, we set up an allocation task. Subjects were recruited from ZJU and SUFE by means of announcements on electronic bulletin boards, which are commonly used to recruit subjects for economics experiments at these universities. Some subjects were randomly assigned to be allocators. Their task was to divide a sum of money between two individuals. The allocation was determined either by the reported private die-rolling outcome or at the allocator's discretion or free will. While all allocators were ZJU undergraduate students, recipients could be either undergraduate students from ZJU (in-group members) or from SUFE (out-group members).

Subjects played only one role, either that of an allocator or that of a recipient. All subjects received a 10 Chinese Yuan (¥10) show-up fee in addition to their earnings in the experiment. We used a between-person approach with a two-by-two factorial design resulting in the four treatments described below. Each subject participated in only one of the four treatments.

The focal treatment was the allocation via a die-roll between two passive recipients, one of whom was an in-group member while the other was an out-group member. We labeled this treatment DieRoll-Ingroup. Specifically an allocator was instructed to allocate ¥50 between two passive recipients.[2] One of these two recipients was another student subject at the same university (ZJU), while the other was a student subject at a different university (SUFE). The allocator was told truthfully that the two participants were only passive recipients and would not

[2] This was equal to approximately 8 US dollars.

make any allocation decisions. The fact that the two recipients were completely passive in the sense that neither would be an allocator during the experiment is an important design feature to isolate possible "generalized reciprocity" motivations from in-group favoritism. The idea of "generalized reciprocity" (see Yamagishi, 2007 for a review of this literature) suggests that in-group bias may arise from the expectation that favorable actions will be positively reciprocated within one's social group. Such reciprocity is not possible under our design.

To investigate cheating, we adopted the die-under-cup method developed by Fischbacher and Föllmi-Heusi (2013). Specifically, allocators were instructed that the determination of the payoff distribution between the two recipient participants would be achieved as follows: "You will be asked to throw a die privately twice. Your first throw will decide which one of the 6 options will be implemented to determine the payoff distribution of the two recipient participants. The second throw only serves to make sure that the die is working properly. You can see the exact payoffs from the following table. You may of course throw the die more than twice. However, only the first throw counts." Then we presented Table 1 below, which contains the six possible allocation options. Notice, that following Fischbacher and Föllmi-Heusi (2013), allocations to in-group recipients rise monotonically from numbers 1 to 5, while the number 6 is associated with a zero allocation to the in-group member. We retain this feature primarily for consistent comparison with other papers using this methodology, but also because it is easy for subjects to remember that the numbers 1 to 5 are associated with the ¥10 to ¥50 in-group allocations they represent. Also following Fischbacher and Föllmi-Heusi (2013), the available allocations do not include an equal split. This is representative of the many situations in life where equal splits are not possible: for example, one person must be chosen from two equally qualified people for the preferred job, while the other must settle for second best. We conjecture that it is in this kind of challenging decision-making situation that subtle forms of in-group favoritism may have a significant impact even on fair-minded inequity-averse allocators.

Insert Table 1 about here.

In addition to the die-roll treatment, we ran a free-will treatment, designed as a control to observe revealed in-group favoritism. Specifically, in the free-will treatment the procedure and the instructions followed as closely as possible those used in the die-roll treatment. The only difference is that subjects were told to “choose one of the 6 options as your preferred allocation for the two recipient participants.” We label this the FreeWill-Ingroup treatment. Notice that because an equal split is not a permitted option, even those participants whose first choice might be to allocate equally between the in-group and out-group member must choose to favor one or the other. If group favoritism were entirely absent among such participants, we would expect approximately half of them to choose the outcome favoring the in-group member and half to choose the outcome favoring the out-group member.

Besides the die-roll vs. free-will dimension, we also used an in-group vs. self dimension in our experimental design. The resulting two-by-two factorial design permits us to compare cheating behavior when the beneficiary of cheating is one’s in-group member versus cheating behavior when the beneficiary is oneself. Specifically, besides the two treatments in which the allocator was asked to allocate a sum of money between two passive recipients described above, we employed two additional treatments in which the allocator had to allocate a sum of money between him/herself and another student from a different university. These two treatments, which we label as DieRoll-Self and FreeWill-Self are designed as further control treatments, allowing a comparison between favoritism and cheating on behalf of oneself versus favoritism and cheating on behalf of an in-group member. Table 2 summarizes the four treatments: DieRoll-Ingroup, FreeWill-Ingroup, DieRoll-Self, and FreeWill-Self.

Insert Table 2 about here.

After allocators completed their allocation task, they were given a post-experimental survey to obtain demographic information such as age, gender and study major. In addition to the ¥10 show-up fee, the allocator received a fixed payment for the allocation task that was independent of the allocation. The fixed payment was equal to the expected payoff of a recipient, namely ¥25. We did not reveal the actual amount of this fixed payment to the allocator until the allocator completed the allocation task in order to avoid influencing the allocation decisions.

In total, we recruited 900 subjects for this experiment. At ZJU we recruited 540 subjects. Of these, 360 were assigned at random to be allocators in one of the four treatments. The remaining 180 were randomly assigned to be recipients in either the DieRoll-Ingroup or FreeWill-Ingroup treatment. (Recall that the DieRoll-Self and FreeWill-Self treatments do not require in-group recipients because the allocation is made between the allocator him/herself and an out-group recipient.) At SUFE 360 students were recruited as recipients and randomly assigned to one of the four treatments. Of the allocators, 65% were female and the average age was 21. The gender ratio was balanced over the four treatments. We did not collect demographic information from the recipients because no such information was available to influence the allocators and recipients made no decisions themselves, playing a purely passive role.

We conducted three allocator sessions of the DieRoll-Ingroup treatment (90 allocators), four allocator sessions of the FreeWill-Ingroup treatment (90 allocators), three allocator sessions of the DieRoll-Self treatment (90 allocators), and three allocator sessions of the FreeWill-Self treatment (90 allocators). Subsequent to each allocator session, there was a corresponding recipient session at each university (90 recipients in each session at each university). The purpose of the recipient sessions was simply to implement the decisions made by the allocators by distributing the funds allocated to the in-group (at ZJU) and out-group (at SUFE) recipients randomly matched to each allocator. Recipient sessions were run at a different time than allocator sessions. We carefully chose this design to guarantee anonymity and ensure that any observed in-group bias did not emerge for strategic or reputational reasons. Thus, while allocators knew whether recipients were from their own University or another University, they did not know anything else about them.

All sessions were conducted by hand. Payments were made in cash, which was given to subjects in private after the session. In addition to the ¥10 show-up fee, average earnings from the experiment were ¥25 for both allocators and recipients. The experiment took about ten minutes. Students at these universities could earn between ¥15-25 an hour at an on-campus job.

4. BEHAVIORAL HYPOTHESES

4.1 Predictions for Each Treatment

Three of our four treatments are either identical or similar to experimental designs that have been examined previously in the literature. Their purpose is to act as controls for the DieRoll-Ingroup treatment, which is the primary focus of this study. We anticipate that we will replicate earlier studies. For example, many studies using a Dictator game setup have demonstrated that although people will give away some money to an anonymous recipient, they will on average keep more than an equal share of the money for themselves (e.g. Kahneman et al., 1986). Our FreeWill-Self treatment is essentially just such a dictator game with six permitted allocations. It differs from most dictator games in that an equal split is not permitted as illustrated in Table 1. Thus, participants are forced to show bias toward either themselves or others. Since results from many Dictator-game studies show that average allocations favor the self over others even when equal splits are permitted, we anticipate that selected allocations will also favor the self over others in our setup where equal splits are not available.

H1: In the FreeWill-Self treatment, subjects will allocate more than half the money to themselves and less than half to a passive recipient.

The DieRoll-Self treatment is identical to the Externality treatment in Fischbacher and Föllmi-Heusi (2013). We expect to replicate their results showing a substantial amount of cheating to achieve allocations more favorable to the allocator.

H2: In the DieRoll-Self treatment, the distribution of the reported die rolls will result in more money for the allocator than would occur under an unbiased uniform distribution of the six possible allocation outcomes.

The FreeWill-Ingroup treatment resembles a number of recent experiments that investigate whether subjects will favor a member of their in-group over an out-group member using both induced groups (e.g., Chen and Li, 2009, Chen and Chen, 2011) and natural groups (e.g., Ockenfels and Werner, 2014). Ockenfels and Werner (2014), like us, use students from two different universities to represent in-group versus out-group members though their allocation task differs from ours. These papers all find favoritism toward in-group members when allocations are freely selected. We expect the same result.

H3: In the FreeWill-Ingroup treatment, subjects will allocate more money to in-group members than to out-group members.

In the DieRoll-Ingroup treatment, subjects have to resort to cheating to favor their in-groups. Combining the results from Fischbacher and Föllmi-Heusi (2013), who report considerable amounts of cheating behavior to benefit oneself, and the results from the literature on in-group favoritism, we conjecture that people will sometimes cheat to benefit their in-group members to the detriment of out-group members. A primary purpose of this paper is to investigate this hypothesis.

H4: In the DieRoll-Ingroup treatment, the distribution of the reported die rolls will result in more money being allocated to in-group members than would occur under an unbiased uniform distribution of the six possible allocation outcomes.

4.2 Predictions Across Treatments

The FreeWill-Self treatment yields a benchmark of revealed preference for the allocation choice between oneself and another anonymous person. In the absence of cheating, the DieRoll-Self treatment should result in allocations drawn from the uniform distribution. Thus, comparing the results of the DieRoll-Self treatment with those of the FreeWill-Self treatment allows us to investigate whether the necessity of cheating to achieve the preferred allocation will constrain people from making that choice. We hypothesize that the moral burden of cheating will inhibit some people from achieving their preferred allocations, but that some cheating will nonetheless occur as predicted by H2.

H5: The average amount allocated to oneself in the DieRoll-Self treatment will be less than that in the FreeWill-Self treatment. Nonetheless, allocations will still be biased towards oneself as predicted by H2.

An analogous trade-off exists for the allocator in the DieRoll-Ingroup treatment. However, the situations are not exactly the same. On the one hand, it is possible that the allocator may not feel as strongly about favoring an in-group member as s/he feels about favoring him/herself, and may therefore be less inclined to cheat to achieve the favored allocation. On the other hand, an allocator may feel less of a moral burden when engaging in cheating on behalf of a group member than s/he may feel engaging in self-interested cheating on behalf of him/herself. Cheating for other in-group members could be rationalized as loyalty to one's group or generous other-regarding behavior, the virtues of which outweigh the moral mandate from the experimental instructions to use the die exclusively as the allocation device. This could result in more cheating. We conjecture that as in the DieRoll-Self treatment, some cheating will take place. The empirical results will allow a comparison between the constraining effects of having to cheat to achieve one's preferred allocation between oneself and another with the constraining effects of having to cheat to achieve one's preferred allocation between an in-group and out-group member.

H6: The average amount allocated to the in-group person in the DieRoll-Ingroup treatment will be less than that in the FreeWill-Ingroup treatment. Nonetheless, allocations will still be biased towards the in-group member as predicted by H4.

5. RESULTS

Insert Figure 1 and Table 3 about here.

Table 3 summarizes the proportion of subjects under each of the six allocations across the four treatments. Figure 1 provides the corresponding cumulative distribution functions for each treatment. Table 3 also provides two-sided binomial tests on whether these proportions are significantly different from the 1/6 probability of each outcome on a die roll, and Kolmogorov-Smirnov one-sample tests on whether the distributions of allocations are significantly different from the uniform distribution. Table 4 reports the average allocations to self in the two "Self" treatments and the average allocations to one's in-group member in the two In-group treatments.

Statistical tests based on a dummy variable regression with average allocations to self or an in-group member as the dependent variable and the treatments as independent (dummy) variables are also reported.[3]

Insert Table 4 about here.

5.1 Results for each Behavioral Hypothesis

We present the results numbered to correspond with each behavioral hypothesis below.

R1: In the FreeWill-Self treatment, subjects on average allocate more than half the money to themselves and less than half to a passive recipient corroborating H1.

In the FreeWill-Self treatment, only 2.2% of the subjects allocate less money to themselves than to the out-group recipient. The two-sided Binomial tests for each cell and a Kolmogorov-Smirnov one-sample test for the whole distribution confirm that the resulting distribution of allocations is significantly different from the uniform distribution (p = 0.000 for the KS test) that would result if these allocations were instead determined by a die roll. On average, allocators give themselves ¥39.11 out of ¥50.00, which according to a t test is significantly different from the equal allocation of ¥25 (p = 0.000). A non-parametric sign test also yields p = 0.000. It is not surprising that H1 is strongly corroborated, replicating many similar results in the Dictator-game literature. Although almost everyone favors him/herself over the anonymous recipient, 68.9% of the allocators nonetheless give some of the money to the recipient.

R2: In the DieRoll-Self treatment, the distribution of the reported die rolls results in more money for the allocator than would occur under an unbiased uniform distribution of the six possible allocation outcomes, corroborating H2.

The DieRoll-Self treatment, like the "Externality" treatment in Fischbacher and Föllmi-Heusi (2013) indicates that some subjects lie to favor themselves over an anonymous out-group person. The two-sided Binomial tests show that the three allocations that favor the other person

[3] An ordered logit regression, not reported in detail to save space, gives qualitatively identical results. They are available from the authors upon request.

occur significantly less frequently than predicted by the random throw of a die, while one of the allocations favoring the allocator (¥40, ¥10) occurs with significantly greater frequency. A Kolmogorov-Smirnov test confirms that the distribution of allocations is significantly different from the uniform distribution associated with the die roll (p = 0.000), evidence that some subjects lie to obtain more money for themselves. On average, allocators take ¥33.44 for themselves, which according to a t test is significantly different from the equal allocation of ¥25 (p = 0.000). A non-parametric sign test also yields p = 0.000. This corroborates H2, broadly replicating the corresponding treatment in Fischbacher and Föllmi-Heusi (2013).[4]

R3: In the FreeWill-Ingroup treatment, subjects allocate more money to in-group members than to out-group members, corroborating H3. However, the favoritism exhibited is modest relative to that in the FreeWill-Self treatment.

The FreeWill-Ingroup treatment reveals the preferences of allocators permitted to select freely the allocation between an anonymous student from the same university, representing an in-group member, and an anonymous student from a different university, representing an out-group member. In contrast to the FreeWill-Self treatment, the vast majority of freely chosen allocations (81.1%) are as close as possible to equality. However, given the unavailability of an equal choice, these allocations subdivide into 57.8% favoring the in-group member with a (¥30, ¥20) division and 23.3% favoring the out-group member with a (¥20, ¥30) division. Only 2.2% of the allocations show greater favoritism toward the out-group member, while the remaining 16.7% show greater favoritism toward the in-group member. The two-sided Binomial tests show that allocations are significantly different than predicted by the uniform distribution for four of the six allocations. A Kolmogorov-Smirnov test confirms that the distribution of allocations is significantly different than would be predicted by a die roll (p = 0.000). On average, allocators give ¥29.33 to the in-group member. According to a t test, this is significantly different from the

[4] The distribution of allocations was not identical between our DieRoll-Self treatment and Fischbacher and Föllmi-Heusi's (2013) Externality treatment (henceforth FF). The main difference was that while in our treatment, the highest percentage of subjects (36.7%) chose the (¥40, ¥10) allocation (25.6% of subjects chose the corresponding allocation in FF), in FF the highest percentage of subjects (26.9%) took all the money for themselves. In contrast, 23.3% of the subjects in our DieRoll-Self treatment chose this allocation.

equal allocation of ¥25 (p = 0.001), while a non-parametric sign test yields p = 0.000. This supports H3, and is consistent with the literature on in-group favoritism.

The favoritism exhibited is however relatively modest. While in the FreeWill-Self treatment, 97.8% of the outcomes favor the allocator, only 74.4% of the outcomes favor the in-group member in the FreeWill-Ingroup treatment. A Kolmogorov-Smirnov test shows that the two distributions are significantly different from each other (p = 0.000). Moreover, the average amount allocators select for themselves in the FreeWill-Self treatment, ¥39.11, is significantly higher than ¥29.33, the comparable amount selected for the in-group member in the FreeWill-Ingroup treatment according to both a t test (p = 0.002) and a Wilcoxon Rank-Sum test (p = 0.000).

R4: In the DieRoll-Ingroup treatment, the distribution of the reported die rolls results in more money being allocated to in-group members than would occur under an unbiased uniform distribution of the six possible allocation outcomes, corroborating H4.

In the DieRoll-Ingroup treatment, the most popular allocation is (¥40, ¥10) favoring the in-group, selected by 28.9% of the allocators. A two-sided binomial test rejects the null hypothesis that this is consistent with truthfully reported results from the roll of a die. Moreover, a Kolmogorov-Smirnov test indicates that the reported allocations are not consistent with the uniform distribution (p = 0.000), suggesting that some allocators are lying to favor members of the in-group. On average, ¥28.11 is allocated to the in-group member. A t-test indicates that this is significantly different from the equal allocation of ¥25 (p = 0.018), thus supporting H4. A non-parametric sign test yields a marginal p = 0.072.[5]

R5: The average amount allocated to oneself in the DieRoll-Self treatment is less than that in the FreeWill-Self treatment, corroborating H5.

We have seen in R2 that self-interested cheating occurs. However, the moral burden

[5] Like all the tests reported in this study, this is a two-tail test. A one-tailed sign test of the null hypothesis of equal allocations against the one-sided alternative that the allocation is biased toward the in-group member rejects the null with p = 0.036.

associated with cheating does appear to keep some people from lying to achieve the same outcome that is freely chosen under the FreeWill-Self treatment. For example, while 97.8% of the outcomes favor the allocator in the FreeWill-Self treatment, this percentage falls to 74.4% in the DieRoll-Self treatment. Indeed, a Kolmogorov-Smirnov test indicates that the DieRoll-Self outcomes are drawn from a different distribution than the FreeWill-Self outcomes ($p = 0.006$). Moreover, allocators take ¥39.11 in the FreeWill-Self treatment, but only ¥33.44 in the DieRoll-Self treatment. This difference is significant according to both a t test ($p = 0.002$) and Wilcoxon Rank-Sum test ($p = 0.050$), corroborating H5.

R6: The average amount allocated to the in-group person in the DieRoll-Ingroup treatment is not significantly different from that in the FreeWill-Ingroup treatment, contrary to the prediction of H6. However, the distribution of allocations differs significantly between the two treatments.

Was in-group favoritism curtailed when people had to cross a moral boundary to express it? The distribution of allocations in the FreeWill-Ingroup treatment looks visually very different from the comparable DieRoll-Ingroup allocation. In particular, the former distribution is massed in the two most equal allocations, while the latter has much more mass in the distributional tails. The former is consistent with many allocators inclined towards equality having only a weak preference to favor slightly the in-group over the out-group member when the only alternative was reverse favoritism to the same small degree. The latter is consistent with such allocators simply reporting truthfully the allocation determined by the die roll. In any case, a Kolmogorov-Smirnov test reveals that the two distributions are significantly different from one another ($p = 0.000$). Despite the distributional differences, the average allocations to in-group members, ¥29.33 in the FreeWill-Ingroup treatment and ¥28.11 in the DieRoll-Ingroup treatment are not significantly different from each other as indicated by both a t test ($p = 0.510$) and a Wilcoxon Rank-Sum test ($p = 0.866$). Thus, H6, which predicts that on average less would be allocated to the in-group member under the DieRoll-Ingroup compared to the FreeWill-Ingroup treatment is not supported.

This result is broadly consistent with there being two types of participants: a majority type and a minority type. The majority type chooses outcomes as close as possible to equality in the FreeWill-Ingroup treatment, and having no strong bias toward either recipient, reports truthfully in the DieRoll-Ingroup treatment. The minority type exhibits strong, but not complete, favoritism in the FreeWill-Ingroup treatment, selecting (¥40, ¥10). They lie if necessary to select the same allocation in the DieRoll-Ingroup treatment, not curtailed by the necessity of cheating to achieve their favored allocation.[6]

5.3 Supplementary Results

Do people cheat more for themselves or for another person who is an in-group member? On the one hand, we have shown that favoritism toward oneself in the FreeWill-Self treatment is significantly stronger than favoritism toward an in-group member in the FreeWill-Ingroup treatment. On the other hand, we have demonstrated that the moral burden of cheating significantly reduces the average amount allocated to the allocator self, while it has no such effect on allocation to an in-group member. The former suggests that self-interested cheating would be greater than cheating for in-group members, while the latter suggests the opposite. Comparing the DieRoll-Ingroup with the DieRoll-Self treatment, the two-sample Kolmogorov-Smirnov test for equality of distribution functions marginally rejects the null hypothesis that the two distributions are the same ($p = 0.059$). Moreover, in monetary terms allocators gave ¥28.11 on average to in-group members in the DieRoll-Ingroup treatment, while taking ¥33.44 on average for themselves in the DieRoll-Self treatment. The difference is significant according to both a t test ($p = 0.004$) and a Wilcoxon Rank-Sum test ($p = 0.010$). With an in-group member defined as a student at the same university as the allocator, significantly less bias emerges when cheating is necessary to favor an in-group member than when it is necessary to favor oneself.

[6] In the FreeWill-Ingroup treatment, 11.1% of allocators chose the (¥40, ¥10) allocation. In the DieRoll-Ingroup treatment, assume that all allocators assigned this allocation by a die roll of 4 reported truthfully. This would be about 1/6 = 16.67% of the allocators. Now assume that 11.1% of allocators favor this allocation as in the FreeWill-Ingroup treatment. Of those, 1/6 would roll a 4 and have no need to cheat. This would leave another (5/6)(11.1%) = 9.25%. We would then expect a total of 16.67% truthful reporters plus 9.25% cheaters equaling 25.92% of allocators to make this selection. The actual percentage was slightly higher at 28.9%.

Finally, we examined whether there were any gender differences in behavior. Gender was never a significant factor in any of the four treatments as indicated by regression-based *t* tests on gender dummy variables and Kolmogorov-Smirnov tests comparing the allocation distributions for males and females within each treatment. Moreover, there was no significant interaction between gender and any of the treatment dummy variables indicated by regression-based *t* tests.[7]

6. CONCLUSION

This paper demonstrates that even small amounts of in-group favoritism based on a minimal naturally-occurring group manipulation can lead to cheating behavior on behalf of a member of one's own group. Our experiment is distinct from other studies in four important ways. First, many earlier studies vary either the benefit or the cost of dishonest behavior, such as the probability of being detected or the potential punishment if caught (see Alm et al., 1992 for a survey of a large body of work on tax evasion that focuses on such issues). In our experiment, the probability of detection is zero and the potential financial benefit from cheating is fixed, while the focus is on the beneficiary of cheating, who may be either the allocator him/herself or another in-group member. Second, a generalized expectation of reciprocity (Yamagishi and Kiyonari, 2000) can play no role in explaining in-group favoritism or cheating based on such favoritism in our study because no participant was permitted to play more than one role. An allocator was only an allocator. A recipient was only a recipient. Our experimental design thus creates a clean and unambiguous setting for examining cheating based on in-group favoritism arising from a seemingly small and insignificant natural difference in social categorization. Third, our study abstracts from potential social efficiency concerns by holding the overall payoffs constant. Fourth, studies have repeatedly shown that people are more willing to lie when other people also benefit from their dishonest behavior (Conrads et al., 2013; Diekmann, et al., 1997; Gino and Pierce, 2010; Schweitzer and Hsee, 2002; Wiltermuth, 2011). However, it is unclear in

[7] We thank an anonymous referee for suggesting that we check for gender effects. Since there were no significant gender effects, we do not report the detailed statistical results. However, they are available from the authors upon request.

these situations whether cheating is driven by self-interest (since the cheater, who is a member on the team, benefits along with everyone else from the cheating behavior) or concern for others (as the benefit also goes to others in the cheater's group). Our design ensures that the cheater gains no direct material payoff from his/her cheating behavior when allocating money between an in-group and out-group member, thus focusing squarely on cheating for another person in one's naturally occurring social group.

Our study also differs in purpose, methodology, and results from the two other recent studies concerning cheating to benefit an in-group member described in the literature review. Jiang (2014) is concerned with whether cheating is more prevalent in more corrupt countries or more collectivist countries, a research question entirely different from our focus on whether people will cheat to allocate more to an in-group member at the expense of an out-group member. Accordingly, opportunities to cheat in her study do not involve the allocation of resources between in-group and out-group members. Rather such opportunities involve the allocation of resources to a subject in the experiment at the expense of the experimenter's research grant. While we employ a between-person design, Jiang (2014) uses a within-person design, and while we use a one-shot game, Jiang's games are repeated. In particular, her subjects have 20 opportunities to cheat for themselves, followed by 15 opportunities to cheat on behalf of a member of a lab-created in-group, and 15 opportunities to cheat on behalf of a member of another lab-created group administered in the same order for all subjects. While Jiang (2014) finds evidence of all three types of cheating in China, her experiment was not designed to compare free-will allocations between an in-group and out-group member with allocations requiring cheating to favor an in-group over an out-group member. Nor does her paper report on whether there was significantly more cheating during the 15 rounds in which in-group members would be the beneficiaries than during the subsequent 15 rounds in which out-group members would benefit.

Hruschka et al. (2014) were also concerned with comparing cheating across countries as outlined in our literature review. Like us, they used natural groups. While our treatments are all based on a between-person design, each allocator in Hruschka et al. (2014) participates in the two allocation tasks described in the literature review in a counter-balanced within-person framework. The between-person and within-person designs each have advantages and disadvantages. Since one of our primary objectives was to examine the difference between cheating for oneself and cheating for others, we did not want to set up an explicit comparison in the minds of our allocators between favoring him/herself and favoring an in-group member. Such a comparison could produce an anchoring effect encouraging either similar allocations in the two treatments (I cheated for another person so I think I am justified in cheating just as much for myself) or alternatively quite different ones (I cheated a little for my in-group member so I think I should cheat more for myself). Our between-subject design ensures that each allocation decision is made in isolation on its own merits.

An important component of our study was to enable the identification of cheating by comparing it to a comparable treatment in which allocations could be freely chosen. Hruschka et al. (2014) did not use such controls. In their Chinese university experiments, 50.2% of the resource was allocated to oneself in the self-other treatment, while the comparable percentage allocated to an in-group member was 50.1%. Thus, there is no evidence of departure from the impartial allocation rule in either treatment. This result differs from the moderate amount of cheating we observe in favor of oneself and the modest amount of cheating we observe in favor of an in-group member. Most importantly, while we are able to compare favoritism when allocations are freely chosen with favoritism requiring cheating, Hruschka et al. (2014) cannot do so. Thus it is unclear whether the equal division of the resource that they observe on average in both of their treatments reflects a lack of self or in-group favoritism or an unwillingness to break the impartial allocation rule.

Earlier research on in-group favoritism has identified several factors that either enhance or mitigate in-group favoritism such as group status, category salience, and common identity (Mullen et al, 1992). Future research should attempt to empirically explore whether these factors will have such effects when in-group favoritism is manifested in the domain of unethical behavior such as cheating and lying. Another route for future inquiry stems from the fact that, in our design, in-group enhancement was necessarily also out-group discrimination (Brewer, 1999). However, this is not always the case. Thus future research could attempt to orthogonally isolate these two issues and explore which is the driving force behind the observed cheating.

Finally, representing in-groups and out-groups by the university a subject attends does not permit random assignment to groups. Although there appear to be many demographic similarities between the students at ZJU and those at SUFE, there may be differences because students are not randomly assigned to universities. For example, ZJU is a larger university with a greater diversity of majors. All of the decision makers in our study were from ZJU. However, it is possible that when making allocation decisions between a ZJU student and a SUFE student, ZJU allocators were influenced by specific perceptions of what kind of student chooses to attend each university, perceptions that may not be generalizable to other groups and populations. While acknowledging that this is a limitation of our study, we would argue that every group has its own unique identity and every group interaction is *sui generis*. We are not claiming that every natural in-group will exhibit in-group favoritism, nor are we claiming that all groups that do so will therefore cheat on behalf of their fellow in-group members. What we are claiming is that with a minimal natural-group manipulation, we have found modest amounts of in-group favoritism and demonstrated that even such modest amounts of in-group bias can lead to unethical behavior. Now that this has been established, it would be interesting to explore the extent to which stronger manipulations might lead to greater amounts of cheating. For example, cheating to favor one’s in-group might be far more prevalent among mutually competitive or hostile groups, and even make one a hero in times of war or severe inter-group strife.

References:

Akerlof, G., and Kranton, R. (2000). Economics and identity. *Quarterly Journal of Economics*, 115(3): 715–753.

Akerlof, G., and Kranton, R. (2002). Identity and schooling: Some lessons for the economics of education. *Journal of Economic Literature*, 40(4): 1167–1201.

Alm, J., McClelland, G., and Schulze, W. (1992). Why do people pay taxes? *Journal of Public*

Economics, 48(1): 21–38.
Alm, J. and M. McKee. (1998). Extending the lessons of laboratory experiments on tax compliance to managerial and decision economics. *Managerial and Decision Economics*, 19: 259–275.
Association of Certified Fraud Examiners (ACFE) (2014). *Report to the nation on occupational fraud and abuse.* Accessed on April 1, 2016 - http://www.acfe.com/rttn-download-2014.aspx
Bandura, A. (1999). Moral disengagement in the perpetration of inhumanities. *Personality and Social Psychology Review*, 3: 193-209.
Basu, K. (2006). Identity, trust and altruism: Sociological clues to economic development. CAE Working Paper: #06-06.
Becker, G. (1957). *The Economics of Discrimination*. The University of Chicago Press.
Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2): 169–217.
Bernhard, H., Fehr, E. and Fischbacher, U. (2006a). Group affiliation and altruistic norm enforcement. *American Economic Review*, 96(2): 217–221.
Bernhard, H., Fischbacher, U. and Fehr, E. (2006b). Parochial altruism in humans. *Nature*, 442: 912–915.
Brewer, M. (1999). The psychology of prejudice: Ingroup love and outgroup hate? *Journal of Social Issues*, 55 (3): 429–444.
Brewer, M. and Yuki, M. (2007). Culture and social identity. In S. Kitayama and D. Cohen (Eds.), *Handbook of Cultural Psychology*, p. 307–322. New York: The Guilford Press.
Buchan, N., Croson, R. and Dawes, R. (2002). Swift neighbors and persistent strangers: a cross-cultural investigation of trust and reciprocity in social exchange. *American Journal of Sociology*, 108(1): 168–206.
Buchan, N., Johnson, E. and Croson, R. (2006). Let's get personal: an international examination of the influence of communication, culture and social distance on other regarding preferences. *Journal of Economic Behavior and Organization*, 60: 373–398.
Buffett, W. E. (2003). Letter to the shareholders of Berkshire Hathaway Inc. for 2002. Accessed on April 1, 2016 - http://www.berkshirehathaway.com/letters/2002pdf.pdf.
Cadsby, C. B., Song, F., and Tapon, F. (2010). Are you paying your employees to cheat? An experimental investigation. *The B.E. Journal of Economic Analysis & Policy*, 10, 1, Article 35.
Charness, G., and Dufwenberg, M. (2006). Promises and partnership. *Econometrica*, 74: 1579–1601.
Chen, Y., and Li, X. (2009). Group identity and social preferences. *American Economic Review*, 99(1): 431–457.
Chen, R. and Chen, Y. (2011). The potential of social identity for equilibrium selection. *American Economic Review*, 101 (6): 2562–2589.
Chugh, D., Bazerman, M., and Banaji, M. (2005). Bounded ethicality as a psychological barrier to recognizing conflicts of interest. In D. Moore, D. Cain, G. Loewenstein, and M. Bazerman (Eds.), *Conflict of interest: Challenges and solutions in business,law, medicine, and public policy*, p. 74–95. New York: Cambridge University Press.
Croson, R., Marks, M., and Snyder, J. (2008). Groups work for women: Gender and group identity in the provision of public goods. *Negotiation Journal*, 24 (4): 411–427.
Conrads, J., Irlenbusch, B., Rilke, R., and Walkowitz, G. (2013). Lying and team incentives.

Journal of Economic Psychology, 34: 1–7.
Diekmann, K., Samuels, S. M., Ross, L., and Bazerman, B. M. (1997). Self interest and fairness in problems of resource allocation: Allocators versus recipients. *Journal of Personality and Social Psychology*, 72: 1061–1074.
Eckel, C., and Grossman, P. (2005). Managing diversity by creating team identity. *Journal of Economic Behavior and Organization*, 58(3): 371–392.
Ehrlich, I. (1973). Participation in illegitimate activities: A theoretical and empirical investigation. *Journal of Political Economy*, 81(3): 521–565.
Erat, S., and Gneezy,U. (2012).White lies. *Management Science*. 58(4): 723–733.
Eriksson, T., Mao, L., Villeval, M. C. (2015). Saving face and group identity. GATE Working Paper 1515.
Fischbacher, U., and Föllmi-Heusi, F. (2013). Lies in disguise: An experimental study on cheating. *Journal of European Economic Association*, 11(3): 525–547.
Gino, F., and Ariely, D. (2012). The dark side of creativity: Original thinkers can be more dishonest. *Journal of Personality and Social Psychology*, 102: 445–459.
Gino, F., and Pierce, L. (2010). Lying to level the playing field: Why people may dishonestly help or hurt others to create equity. *Journal of Business Ethics*, 95: 89–103.
Gneezy, Uri (2005). Deception: The role of consequences. *American Economic Review*, 95(1): 384–394.
Harris, D., Herrmann, B., Kontoleon, A., and Newton, J. (2015). Is it a norm to favour your own group? *Experimental Economics*, 18: 491–521.
Herman, T. (2005). Study Suggests tax cheating is on the rise; Most detailed survey in 15 years finds $250 billion-plus gap; Ramping up audits on wealthy. *The Wall Street Journal*, (March 30), D1.
Houser, D., Vetter, S., and Winter, J. (2012). Fairness and cheating. *European Economic Review*, 56(8): 1645–1655.
Hogg, M., and Abrams, D. (1993). Towards a single-process uncertainty-reduction model of social motivation in groups. In *Group Motivation: Social Psychological Perspective*, M. A. Hogg and D. Abrams (Eds.), p. 173–190. London: Harvester Wheasheaf.
Hruschka, D., Efferson, C., Jiang, T., Falletta-Cowden, A., Sigurdsson, S., McNamara, R., Sands, M., Munira, S., Slingerland, E., and Heinrich, J. (2014). Impartial institutions, pathogen stress and the expanding social network. *Human Nature*, 25: 567–579.
Jensen, M. C. (2001). Corporate budgeting is broken – Let's fix it. *Harvard Business Review*, 79: 94–101.
Jensen, M. C. (2003). Paying people to lie: The truth about the budgeting process. *European Financial Management*, 9: 379–406.
Jiang, T. (2014). Other-regarding preferences and other-regarding cheating – Experimental evidence from China, Italy, Japan and the Netherlands. SSRN Working Paper.
Jones, T. M. (1991). Ethical decision making by individuals in organizations: An issue-contingent model. *Academy of Management Review*, 16(2): 366–395.
Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1986). Fairness and the assumptions of Economics. *The Journal of Business*, 59 (4): 285–300.
Klor, E. and Shayo, M. (2010). Social identity and preferences over redistribution. *Journal of Public Economics*, 94: 269–278.
Mazar, N., Amir, O., and Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research*, 45(6): 633–644.

McCabe, D. L. and Trevino, L., K. (1993). Academic dishonesty: Honor codes and other contextual influences. *Journal of Higher Education*, 64 (5): 522–538.
McCabe, D. L. and Trevino, L., K. (1997). Individual and contextual influences on academic dishonesty: A multi-campus investigation. *Research in Higher Education*, 38 (3): 379–396.
McLeish, K., and Oxoby, R. (2011). Social interactions and the salience of social identity. *Journal of Economic Psychology*, 32: 172–178.
Mullen, B., Brown, R., and Smith, C. (1992). Ingroup bias as a function of salience, relevance, and status: An integration. *European Journal of Social Psychology*, 22(2): 103–122.
Nagin, D. S., Rebitzer, J., Sanders, S., and Taylor, L. (2002). Monitoring, motivation and management: The determinants of opportunistic behavior in a field experiment. *American Economic Review*, 92(4): 850–872.
Ockenfels, A., and Werner, P. (2014). Beliefs and ingroup favoritism. *Journal of Economic Behavior and Organization,* 108: 453–462.
Reynolds, S. J. (2006). Moral awareness and ethical predispositions: Investigating the role of individual differences in the recognition of moral issues. *Journal of Applied Psychology,* 91(1): 233–243.
Ruffle, B. J. and Sosis, R. (2006) Cooperation and the in-group-out-group bias: A field test on Israeli Kibbutz members and city residents. *Journal of Economic Behavior and Organization,* 60: 147–163.
Schweitzer, M. E. and Hsee, C. K. (2002). Stretching the truth: Elastic justification and motivated communication of uncertain information. *Journal of Risk and Uncertainty,* 25: 185–201.
Schweitzer, M. E., Ordóñez, L., and Douma, B. (2004). Goal-setting as a motivator of unethical behavior. *Academy of Management Journal*, 47(3): 422–433.
Shalvi, S., Dana, J., Handgraaf, M., and DeDreu, C. (2011). Justified ethicality: Observing desired counterfactuals modifies ethical perceptions and behavior. *Organizational Behavior and Human Decision Processes*, 115: 181–190.
Shalvi, S., Eldar.,O., and Bereby Meyer, Y. (2012). Honesty requires time (and lack of justifications). *Psychological Science*, 23: 1264–1270.
Shalvi, S., and Leiser, D. (2013). Moral firmness. *Journal of Economic Behavior and Organization,* 93: 400–407.
Solow, J. and Kirkwood, N. (2002). Group identity and gender in public goods experiments. *Journal of Economic Behavior and Organization*, 48(4): 403–412.
Song, F., Cadsby, C. B., and Bi, Y. (2012). Trust, reciprocity and *Guanxi* in China: An experimental investigation. *Management and Organization Review*, 8(2): 397–421.
Song, F., and Zhong, C. B. (2015). You scratch his back, he scratches mine, and I’ll scratch yours: Moral hazard in direct and indirect reciprocal agency. *Journal of Economic Behavior and Organization*, 112: 98–111.
Tajfel, H., and Turner, J. (1979). An integrative theory of intergroup conflict. In *The Social Psychology of Intergroup Relations,* W. G. Austin and S. Worchel (Eds.), p. 33–47. Monterey: Brooks-Cole.
Tajfel, H., and Turner, J. (1985). The social identity theory of intergroup behavior. In *The Psychology of Intergroup Relations,* W. G. Austin and S. Worchel (Eds.), p. 7–24. Chicago: Nelson-Hall.
Tajfel, H., Billig, M., Bundy, R., and Flament, C.L. (1971). Social categorization and inter-group behavior. *European Journal of Social Psychology,* 1(2): 149–178.

Triandis, H., Bontempo, R., Villareal, M., Asai, M., and Lucca, N. (1988). Individualism and collectivism: cross-cultural perspectives on self–ingroup relationships. *Journal of Personality and Social Psychology*, 54(2): 323–338.

Wiltermuth, S. S. (2011). Cheating more when the spoils are split. *Organizational Behavior and Human Decision Processes*, 115: 157–168.

Yamagishi, T. (2007). The social exchange heuristic: A psychological mechanism that makes a system of generalized exchange self-sustaining. In Radford M., Ohnuma, S. Yamagishi, T. (Eds.), *Cultural and Ecological Foundations of the Mind*, p. 11–37. Japan: Hokkaido University Press.

Yamagishi T. and Kiyonari T. (2000). The group as the container of generalized reciprocity. *Social Psychology Quarterly*, 63(2): 161–197.

Web References:

http://www.biaozhun007.com Accessed on April 1, 2016.

http://www.cuaa.net Accessed on April 1, 2016.

Table 1 Allocations Corresponding to Die-Roll Outcomes

Die-Roll Outcome	1	2	3	4	5	6
Payoff of the participant from your University	¥10	¥20	¥30	¥40	¥50	¥0
Payoff of the participant from another University	¥40	¥30	¥20	¥10	¥0	¥50

Table 2 Experimental Manipulations

Allocation Rule / Recipients	Die-Rolling	Free-Will
An in-group member and an out-group member	DieRoll-Ingroup	FreeWill-Ingroup
Self and an out-group member	DieRoll-Self	FreeWill-Self

Figure 1 Percentage of Subjects in Each of the Six Outcomes: Cumulative Distribution Functions.

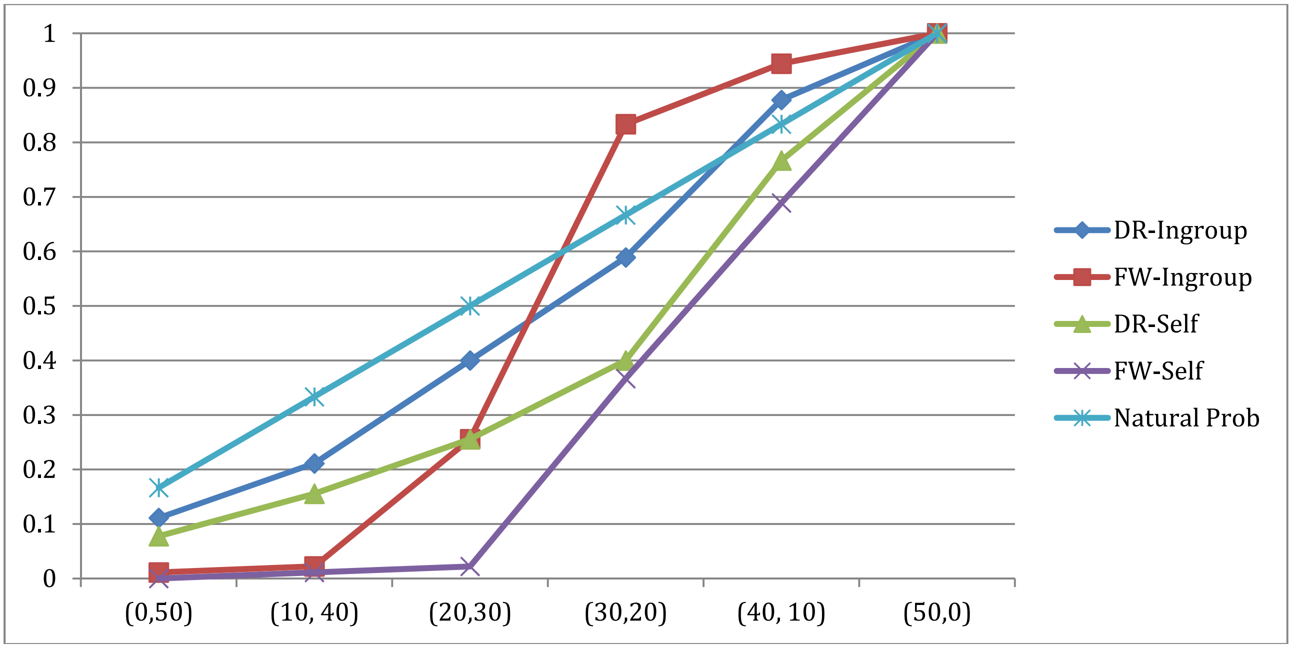

Table 3 Percentage of Subjects in Each of the Six Outcomes and Two-sided Binomial Tests

Allocation (In-, Out-group) or (Self, Other)	FW-Self Treatment (n=90)	DR-Self Treatment (n=90)	FW-Ingroup Treatment (n=90)	DR-Ingroup Treatment (n=90)
(¥0, ¥50)	0%***	7.8%**	1.1%***	11.1%
(¥10, ¥40)	1.1%***	7.8%**	1.1%***	10%*
(¥20, ¥30)	1.1%***	10%*	23.3%	18.9%
(¥30, ¥20)	34.4%***	14.4%	57.8%***	18.9%
(¥40, ¥10)	32.2%***	36.7%**	11.1%	28.9%***
(¥50, ¥0)	31.1%***	23.3%	5.6%***	12.2%
Kolmogorov-Smirnov One-sample Test	0.000	0.000	0.000	0.000

Note: *, ** and *** denote *p*-value at 10%, 5%, and 1% of the one-sided binomial test that the observed frequency is smaller (larger) than 16.7%.

Table 4 Average Allocations to Ingroup/Self and Associated Parametric Tests

Panel A: Average Allocations to Self/Ingroup across Treatments

Recipients \ Allocation Rule	Die-Roll	Free-Will
In-Outgroup	¥28.11	¥29.33
Self-Other (Outgroup member)	¥33.44	¥39.11

Panel B: Dummy Variable Regression on Average Allocations

	Coef.	Std. Error	*p*-value
Constant (FreeWill-Ingroup Treatment)	29.33	1.31	0.000
Die-Roll (DieRoll=1, FreeWill=0)	-1.22	1.85	0.510
Self (Self=1, Ingroup=0)	9.78	1.85	0.000
Interaction	-4.44	2.62	0.091

Panel C: Two-Tailed Hypothesis Tests (*t* tests and non-parametric tests as indicated)

H1: 39.11 > 25, t = 10.76, p = 0.000.	Sign Test: p = 0.000.
H2: 33.44 > 25, t = 6.44, p = 0.000.	Sign Test: p = 0.000.
H3: 29.33 > 25, t = 3.30, p = 0.001.	Sign Test: p = 0.000.
H4: 28.11 > 25, t = 2.37, p = 0.018.	Sign Test: p = 0.072.
H5: 39.11 > 33.44, t = 3.06, p = 0.002.	Wilcoxon Rank-Sum Test: p = 0.050.
H6: 29.33 = 28.11, t = 0.66, p = 0.510.	Wilcoxon Rank-Sum Test: p = 0.866.
39.11 > 29.33, t = 5.27, p = 0.002.	Wilcoxon Rank-Sum Test: p = 0.000.
33.44 > 28.11, t = 2.88, p = 0.004.	Wilcoxon Rank-Sum Test: p = 0.010.