UNIVERSITÄT PADERBORN
Die Universität der Informationsgesellschaft

FACULTY OF BUSINESS ADMINISTRATION AND ECONOMICS

Working Paper Series

Working Paper No. 2018-05

Do children stop cheating when someone else reaps the benefits? A lab in the field experiment *

Julia Kramer, Silvia Lübbecke, and Nina Lucia Stephan

January 2019

* This study is a revised version of the originally named draft "Cheating for my or for your benefit? A field experiment with children"published in this series in February 2018.

WiWi

Do children stop cheating when someone else reaps the benefits? A lab in the field experiment

Julia Kramer*

Silvia Lübbecke†

Nina Lucia Stephan‡

January 24, 2019

Abstract

This paper describes a lab in the field experiment with 512 children, investigating whether cheating depends on who receives the benefits from cheating. In the experiment, children privately toss a fair coin with a 50% chance of winning a prize and subsequently report the outcome. Children are partnered with another child. In the treatment group, the prize in case of winning is received by the passive partner, instead of the child tossing the coin. We find that the share of prize winning outcomes is significantly reduced when the prize is received by the partner. Results are strongest when the partner is a stranger. In line with existing literature, this effect is weaker and eventually disappears among older children. We argue that part of the effect is explained by children cheating to prevent the partner from receiving the prize, and hence, engaging in downward cheating.

Keywords: cheating, dishonesty, altruism, children, lab-in-the-field experiment

JEL-Codes: C93; D63; D64; Z13

*julia.kramer@upb.de

†silvia.luebbecke@upb.de

‡nina.stephan@upb.de

1 Introduction

Honesty is an essential ingredient of economic activity. Therefore, economists have long been interested in studying cases of and motives behind dishonest behavior in many areas, such as tax compliance (Andreoni et al., 1998), job applications (Stigler, 1962), honesty at the work place (Tyler and Blader, 2005; Nonis and Swift, 2001), in education (Jacob and Levitt, 2003; Whitley, 1998), in research (McCabe et al., 2001; McCabe and Trevino, 1997) or in politics (Nye, 1967). The present paper contributes to existing literature concerned with the factors driving individuals to behave dishonestly.

Factors driving dishonesty do not only include individual profits and potentially the risk of being detected, but also describe the existence of psychological costs of lying (Abeler et al., 2014; Fischbacher and Föllmi-Heusi, 2013; Thielmann and Hilbig, 2018). It is widely accepted that individuals are motivated to behave honestly to preserve a positive image, not only towards the outside world but also towards themselves (e.g. Barkan et al., 2015; Mazar et al., 2008; Pittarello et al., 2015). Hence, individuals look for ways to excuse their dishonest behavior. The availability of a justification is thus another driving factor underlying dishonest behavior.

One possible justification is to describe own dishonesty as a white lie (Erat and Gneezy, 2012), that is a lie that (also) benefits someone else, more generally referred to as self-serving altruism (Irlenbusch and Villeval, 2015; Shalvi et al., 2015; Klein et al., 2017). For example, Gneezy (2005) or Conrads et al. (2013) show that dishonesty increases when a lie also benefits someone else. Other studies find that even in the absence of individual profits from cheating, individuals engage in cheating to benefit others (e.g. Gino et al., 2013; Wiltermuth, 2011). Klein et al. (2017) compare how the willingness to cheat depends on whether the profit is received by the decision maker, partly received by someone else, or solely received by someone else. While they find that benefits through self-profit outweigh psychological costs of social justifiability, they also find that when the profit is received solely by someone else, cheating rates are not significantly different from those when the profit is received by the decision maker alone.

Even though there is a large literature in economics as well as other social sciences on both, the reasons for and impacts of dishonesty, little is known about how dishonesty develops in childhood and adolescence. We address this by studying decisions made by young children and adolescents. More precisely, we investigate whether the motive to altruistically cheat, despite the absence of individual profits from cheating, already exists among children, and thus very early in life, just as Klein et al. (2017) show it to be the case among adults. In order to do so, we test whether there is more or less cheating, when individual profits from cheating are taken away, and instead the benefits from cheating go to another child.

In the present paper, we describe a lab in the field experiment inspired by Bucciol and Piovesan (2011) who study cheating behavior of children aged 5 to 15 years in a lab in the field experiment. Just as in their study, we ask children to flip a coin with a 50% chance of winning a prize. The coin is flipped in private and the outcome is only observed by the children themselves, hence they can either obtain the prize by winning or by falsely reporting the winning outcome. Bucciol and Piovesan find that a significant share of children does cheat, especially children above the age of eight. Just as in the present study, they investigate how choices depend on age. We take the investigation a step

further by not only studying cheating out of selfish motives. Instead, our central question is how decisions change when cheating does not benefit the cheater himself, but benefits someone else instead, and may thus happen out of other motives. Moreover, we provide 256 independent observations from 521 children aged 3 to 16, including observations from siblings, friends, as well as strangers. This rich dataset enables us to describe how the answer to our central question depends on both, age, and the relationship to the other child.

Our central result is that while a significant share of young children cheat in order to obtain an individual profit for themselves, there is significantly less beneficial cheating when the profit is obtained by someone else instead. Cheating to benefit someone else is relatively larger when the other person is known. Our central result is not only dependent on the relationship to the other child, but also dependent on the age. The result of beneficial cheating being reduced when the profit is obtained by somebody else is strongest among very young children. This finding suggests that whereas selfish motives may be strongest in this subgroup, altruistic motives may not yet have developed among the very young. In contrast, older children make less of a difference between obtaining a profit for themselves, versus obtaining it for someone else, which is in line with literature on adult cheating behavior (Klein et al., 2017).

The remainder of the paper is organized as follows: In Section 2, we formulate expectations regarding our results, based on existing related literature. Section 3 describes the experiment in more detail. Section 4 describes the implementation and participants. Section 5 continues with a summary of the results and their analysis. Section 6 concludes.

2 Theoretical Considerations

Currently, dishonesty is a broadly discussed topic. A recent overview is provided by Rosenbaum et al. (2014), who review 63 economic and psychological experiments on cheating. While most experimental economic literature investigates adults' lying behavior (e.g. Fischbacher and Föllmi-Heusi, 2013; Abeler et al., 2014; Conrads and Lotz, 2015; Houser et al., 2012; Jiang, 2013), very little research has been done to investigate the existence of dishonesty among children, and how dishonest behavior develops with age. Chytilova and Korbel (2014) find that in when deciding in groups, children cheat more than when deciding individually. Glätzle-Rützler and Lergetporer (2015) and Maggian and Villeval (2016) find that older children lie less than younger children, as they are more aware of the harm caused by lies.

As early as with the age of 3, children start understanding what is socially considered right and wrong (Buchmann et al., 2009; Killen and Smetana, 2008). However, according to Nunner-Winkler and Sodian (1988), while they start realizing moral concepts at very young ages of 3 or 4 years, learning to translate this knowledge into actions takes more life experience. In our experiment, the simple decision to cheat in order to gain a profit for oneself becomes more complex in the treatment group, where someone else receives the benefit from cheating. Very young children may not yet have developed motives for cheating other than egoism. This motive disappears, when cheating only benefits a partner. Older children, however, understand that the decision is now also about being altruistic. Hence, additionally to investigating the overall treatment effect, we are also interested in describing in which way the answer depends on the age of the children. We expect to

find that very young children do cheat for themselves, but rather not when cheating solely benefits another child.

Lately, some studies on adults' cheating behavior describe cheating that increases others' payoff. For example, just as in the present paper, Houser et al. (2016) are interested in the question whether cheating changes when it benefits another person, rather than the decision maker himself. They observe that parents are more likely to cheat for their children's benefit than for their own benefit. The present study extends on this by investigating whether this preference for altruistic children is already present during childhood, or whether it develops as people age.

Fehr et al. (2013) study children's and adolescents' behavior, but focus on other-regarding preferences, rather than cheating to benefit others. They find that while choices motivated by spite and egalitarianism decrease with age, choices motivated by altruism increase. In line with this, we expect older children to be more motivated to obtain a profit that only benefits someone else, rather than themselves.

Moreover, several existing studies provide evidence for some individuals engaging in spiteful behavior (e.g. Kimbrough and Reiss, 2012). In other words, some individuals have a preference for making choices that harm, rather than benefit others. Hence, when analyzing cheating in favor of others, we should consider the possibility that some individuals cheat in order to avoid producing a profit for others.

Further existing literature on dishonesty suggests that the decision to benefit someone else is likely also dependent on the relationship of the decision maker to this other person (Belot and Van de Ven, 2011; Chen and Li, 2009). Neuroscientific research offers the explanation that providing to a friend, rather than a stranger, produces more feelings of joy (Fareri et al., 2012). Therefore, we will further specify the results by considering whether the partner being a stranger versus a known child, or even a sibling, changes the results. We expect that when the two children are related in some way, it is easier for the children to understand the positive benefit of cheating in favor of this other child. Hence, they are more likely to cheat in favor of the other child, or in other words, to behave altruistically.

While gender effects are not in the focus of this study, evidence from previous studies shows that female adults (Eckel and Grossman, 1998; Andreoni and Vesterlund, 2001), female adolescents (Dreber et al., 2014), and female children (Houser and Schunk, 2009) tend to behave more generously. From this, we expect that female subjects in our experiment may also behave more generously towards others, and hence to be more likely to cheat in favor of others.

3 Experimental Design

In order to answer our central research question, as well as whether the results depend on age, relationship to the other child, or gender, it is necessary to attract a heterogeneous group of participants. In order to obtain this variation in characteristics, the present experiment is conducted as a lab in the field experiment. In this section, we describe the design of our experiment. First, we describe the setup in the surrounding environment. Second, we explain the sequence of the experiment itself.

3.1 Surrounding Field Environment

We conduct the experiment during the open days of a large museum. Specifically, the experimental area is located within the large entrance hall of the museum. The open day is visited by school classes, kindergarten groups and families, and thus provides a heterogeneous pool of potential participants.

We take several measures to advertise our experiment and attract participants outside the experimental area. Doing so outside the experimental area also allows children to get to know the experimenters, before deciding to actually follow them into the experimental area. First, we attract children by asking for their 'help' testing a die. We invite children to throw the die and mark the outcome on a big flip chart, thereby collecting statistics to show the frequency of the different outcomes (Appendix Figures 4 and 5). Second, we install a poster wall with pictures and information on random devices for the passer-by (Appendix Figure 6). Both attractions are also meant as an entertainment for the siblings, parents, or teachers, who are waiting for their children to finish participating in the experiment.

The experiment is framed such that participating in it appears to the children as playing a game of fortune, as it is all about throwing dice or coins and winning something. Moreover, the experiment is just one out of many stands that are exclusively set up for the occasion of the open day. This ensures that children behave rather normally, and do not make choices different from how they would usually make them.

In order to avoid participants from being distracted or feeling observed while participating in the actual experiment, the experimental area is surrounded by solid walls (e.g. Appendix Figure 5). There is one single entry and exit to the experimental area, so children cannot get lost. Additional walls inside the experimental area are used to guide the children through the different steps of the experiment, as explained in Section 3.2.

Children who decide to participate are asked to wait in a single line in front of the entrance to the experimental area. Upon arrival at the entrance to the experimental area, every other pair is either sent to one side, to participate in the control group, or to the other side, in order to participate in the treatment group. Hence, this is a between-subject design. In order to increase comparability, both control and treatment group sessions take place parallel to each other throughout the experiment. When entire school classes or kindergarten groups come to participate, they are thus split up into both control and treatment group, thereby avoiding randomization issues. Hence, inside the experimental area, two pairs of children can participate in the experiment at the same time. Within the experimental area, the control and treatment group areas are separated from each other with the help of additional walls.

3.2 Sequence

Children participate in pairs of one 'active' and one 'passive' child. We assign roles randomly by letting one child role a two-colored die, and letting the other child bet on the winning outcome, as illustrated in Appendix Figure 7. The winner becomes the active child. After assigning the roles, both children accompany the experimenter into the experimental area to Position 1, as shown in Figure 1.

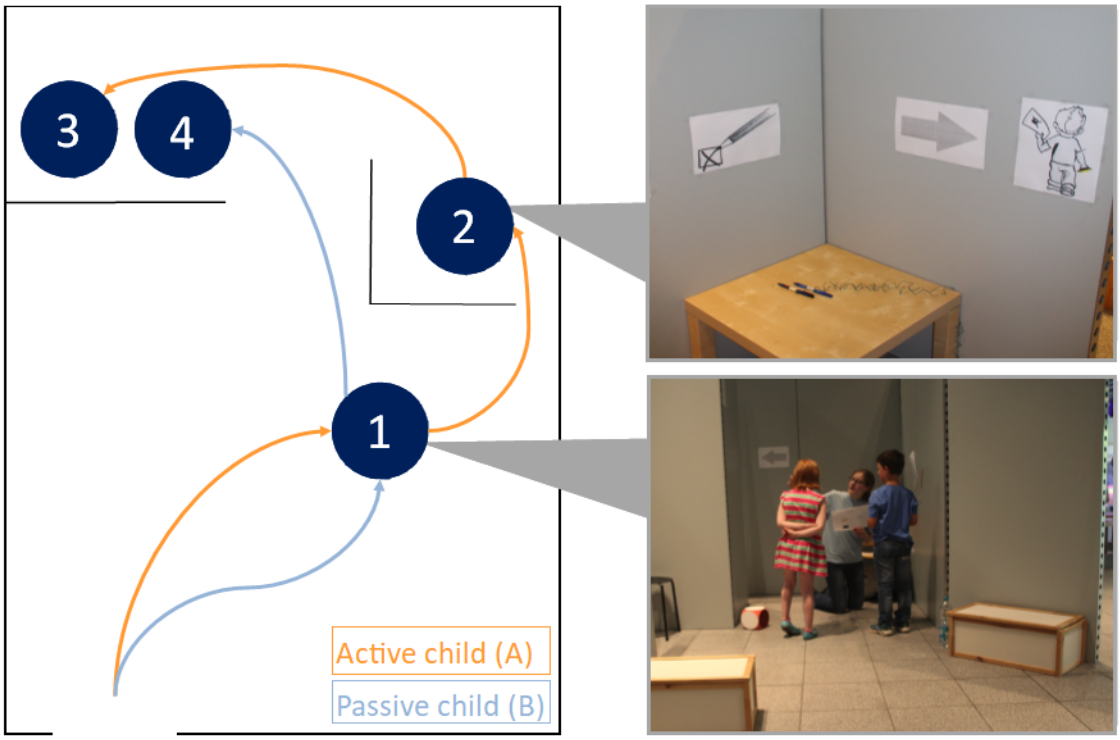

Figure 1: Setup of the experiment with positions of participating children

At Position 1, the experimenter explains the entire set of instructions to both children, carefully ensuring that both understand them fully. Pictures on the walls of the experimental area are used to support the instructions (Figure 1).

In order to avoid asymmetric information between the participants joining the experiment at different times of the day, due to early participants showing off their payoffs to future participants, we take another simple measure. All participants are shown the potential payoffs upfront, as part of being instructed to participate in the experiment. The payoff consists of a fixed payoff in form of a small die. Additionally, one child from each pair of children can win a prize. The prize is a rainbow color pencil, displayed in Appendix Figure 8.

The active child is asked to enter a private corner (Position 2), where no one can observe her actions. The experimenter makes sure that she knows she will not be observed by anyone while carrying out her task in the private corner. She is given a fair yellow-and-black coin, along with instructions to toss the coin in private. She is asked to report the outcome by a mark on the report sheet (Figure 2) before leaving the private corner. Hence, the active child has the possibility to report any color, also the one different from the actual outcome, without any risk of being detected. Next, she hands the report sheet over to a second experimenter, waiting at the exit of the private corner (Position 4).[1]

[1]Within the private corner, additional graphical instructions, displayed in Appendix Figures 9 and 10, remind the active child of what to do.

Winning the prize has different consequences in the control compared to the treatment group: In case of a win in the control group, the prize is obtained by the active child. In the treatment group, the prize is obtained by the passive child instead. Hence, in the treatment group the active child cannot influence her own payoff, but only that of the other child.

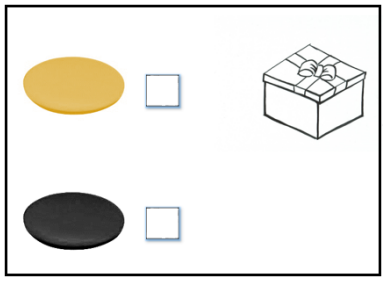

Figure 2: Report sheet, on which the active child is asked to mark yellow or black in the private corner

While the active child is carrying out the main task, the passive child receives a simple assistant task. This is important in order to ensure that both children take the experiment serious, and are not bored or confused. Thus, we use this assistant task to integrate and involve the passive child without impacting or disturbing the active child's behavior. This assistant task consists in walking around the private corner and bringing an empty form to the second experimenter, waiting at Position 4. Hence, both children meet again, once the active child finishes her task (moving from Position 3 to 4).[2]

Next, the second experimenter, after receiving the report sheet and with both children present, evaluates the outcome reported on the sheet. In case of a yellow report in the control group, only the active child is then awarded the rainbow colored pencil. In case of a yellow report in the treatment group, only the passive child is then awarded the rainbow color pencil. In case of a black report, in both the control and treatment group, no one receives the prize. The possible choices and resulting payoffs are displayed in form of a game tree in Appendix Figure 11.

Finally, the experimenter asks a set of control questions and demographic questions printed on the form brought by the passive child. The latter enables analyzing gender- age- and relationship specific treatment effects. The experimenter marks the answers on the form. He awards both children with a small die, as well as a stamp on their hand. Giving the children a stamp mark comes with the advantage that the experimenters outside the experimental area can make sure that no child participates in the experiment more than once.

4 Implementation and Participants

The experiment was carried out in Paderborn, Germany, in June 2015. A total of 512 children participated in the experiment over the course of three consecutive days. As

[2]In most cases, the active child's task did not take longer than a minute, such that there was no long waiting time for the passive child.

children participated in pairs, the dataset contains 256 independent observations. Two children left the experiment without reporting a coin flip outcome and before answering the questionnaire, which is why we exclude these two observations from the dataset for any of the following calculations. Demographic characteristics of the remaining sample of 254 active children are displayed in Table 1. We will refer to these 254 participants, out of the total of 512 participants, as our 'subjects'.

	Control Group	**Treatment Group**	
Observations (254 subjects)	120	134	
Age (min 3, max 15)	7.31	8.10	**
Female (%)	51.26	42.54	
Paired with female (%)	59.66	48.87	*
Partner known (%)	67.22	74.63	
Paired with sibling (%)	24.58	14.93	*

Table 1: differences between control and treatment group (t-test); level of significance denoted * p<0.1, ** p<0.05, *** p<0.01

Running the experiment in the field allowed us to gather observations not only from pairs of strangers, but also from pairs of friends and even siblings. Comparing demographic characteristics of the subjects using a t-test reveals that subjects in the treatment group are significantly older (p=0.01), and weakly significantly less often paired with a sibling (p=0.05) or with a female partner (p=0.09). When evaluating the data, we will thus control for these three variables to make sure that potential differences between the control and treatment group are not only driven by these differences. Regarding the gender of the active children, or the share of children who were paired with a child already known to them, there are no notable differences between the treatment and control group. Here, randomization was successful. Despite a few imperfections in the randomization, we still obtained a rich dataset with the desired heterogeneity in participant characteristics. There are sufficient observations of each type of participant in both the control and the treatment group to allow testing for type-specific effects.

5 Results and Discussion

First, we analyze whether we find evidence for cheating in the overall sample within the control and treatment group, and compare the two. Next, we go into more detail and describe how results depend on the relationship between the children, as well as on age.

5.1 Central Results

Out of the 134 subjects in the treatment group, 64 subjects, a share of 47.76%, reported the prize winning outcome "yellow". In contrast to that, 81 out of 120 subjects in the control group, a share of 67.50%, reported the prize winning outcome. According to a two-sided binomial test, the share of reported prize winning outcomes in the treatment group is not

significantly different from 50%. In contrast, within the control group it is significantly different from 50% ($p<0.05$), which means that a significant share of subjects in the control group is engaging in cheating in order to obtain the prize.

Comparing the "yellow" reports of the treatment group to those of the control group using a Chi-Square test shows that yellow reports are significantly more likely in the control group ($p<0.01$). This difference is illustrated in Figure 3, and leads to our first central result:

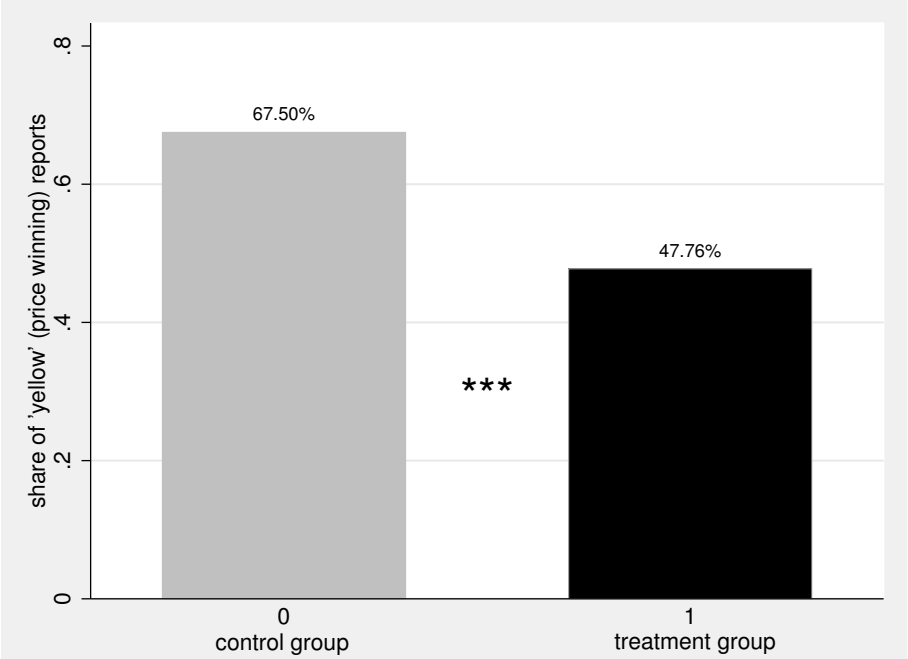

Figure 3: Treatment effect: significantly different shares of subjects report the prize winning outcome of a coin flip (*** $p<0.01$)

Result 1. *Subjects are more likely to report the price winning outcome of a coin flip when the prize benefits themselves, compared to when the prize benefits another child instead.*

We interpret this result with great care, when it comes to drawing conclusions regarding cheating behavior. Cheating, in this case, can go in two directions. First, subjects can falsely report "yellow" in order to win a prize, either for themselves or for someone else. Second, subjects can falsely report "black" in order to avoid winning a prize themselves, or to avoid that someone else receives the prize. In existing literature, the second type of cheating is referred to as downward lying or downward cheating (e.g. Houser et al., 2012; Abeler et al., 2014; Utikal and Fischbacher, 2013). When both types of cheating behavior are present, upward and downward cheating cancel each other out, either fully or partly. Therefore, we cannot conclude that subjects in the treatment group are not cheating, even though the prize winning outcome is not reported in more than 50% of the cases. However, we can conclude that there are relatively more cases of upward cheating, compared to downward cheating, in the control group, while this is not the case in the treatment group. Hence, comparing types of cheating across groups, the share of upward to downward cheating is higher in the control, rather than the treatment group.

5.2 Further Results

To obtain more detailed results and test for type-specific effects, we run Probit regressions,[3] with the prize winning outcome as the dependent variable. Marginal effects are reported in Table 2.

Overall, Probit estimates confirm Result 1. Reporting the prize winning outcome is significantly ($p<0.01$) less likely in the treatment group, where the prize benefits the partners, rather than the subjects themselves. While the size of this effect varies between 19.4 percentage points without controls and 17.6 percentage points with controls, it is robust to controlling for differences between the control and treatment group described in Section 4, regarding age, being paired with a sibling, and being paired with a female (see column 2).

Next, we include interaction terms to test whether the effect differs, depending on whether the subject knows his or her partner, as well as depending on the subject's age, and consider estimates in Table 2 Columns 3 and 4. Testing for gender-specific treatment effects did not yield significant results.[4]

	(1)	(2)	(3)	(4)
treatment	-0.194***	-0.176***	-0.348***	-0.559***
	(0.06)	(0.06)	(0.11)	(0.19)
age_subject		-0.011	-0.010	-0.041**
		(0.01)	(0.01)	(0.02)
female_partner		0.097	0.105*	0.101*
		(0.06)	(0.06)	(0.06)
siblings		-0.042		-0.043
		(0.08)		(0.08)
known			-0.064	
			(0.10)	
known_treat			0.249*	
			(0.13)	
age_treat				0.050**
				(0.02)
N	254	250	251	250

Table 2: Probit marginal effects with reporting "yellow" as dependent variable. Control variables include the subject's age, the partner's gender and controls for whether the subject and partner are siblings or known to each other. We also control for relationship-specific and age-specific treatment effects by including interactions of the variables known and age with the treatment dummy. * p<0.1, ** p<0.05, *** p<0.01, (se)

First, we consider regression estimates in Column 3. Intuitively, the effect of the partner being a sibling is intuitively stronger than the partner only being someone known.

[3]The number of observations used for the different regressions varies due to a few missing questionnaire entries from some children who were not patient enough to answer all questions.

[4]Further regression results, including interaction terms to test for gender-specific effects, are reported in the Appendix Table 3.

However, there are relatively few siblings in the sample. Moreover, in real live, interactions with known persons or strangers are both more frequent than with siblings and hence, we drop the control variable siblings in favor of the control variable known. In line with the previous Result 1, we find a highly significant ($p<0.01$) baseline treatment effect, also when testing for relationship-specific effects. Moreover, it is shown that in the control group whether the passive child is known to the subject does not seem relevant for the decision to report the prize winning outcome. Intuitively, this makes sense since in the control group, partners are not affected by the outcome of the coin flip. In contrast, whether the partner is known or not becomes relevant in the treatment group, where winning benefits the partner: The positive estimate of the interaction between treatment and known (0.249, $p<0.1$) shows that the treatment effect is reduced when children know each other. In other words, the strongly negative treatment effect is driven by pairs of children who are strangers. The effect is smaller when the children know each other. Thus, subjects are more likely to report the prize winning outcome when the prize is received by somebody they know, rather than a stranger.

Result 2. *The size of the treatment effect depends on the relationship between the subject and her partner. When they know each other, the effect is reduced.*

Second, regression estimates in Column 4 show that in the control group, where winning benefits the participants themselves, older children are significantly ($p<0.05$) less likely to report the prize winning outcome. This reflects an unintended drawback of the design: Anecdotal evidence from comments by participants suggest that the pencil with rainbow colors, the winning prize, is less attractive to older children. At the same time, estimates again show a highly significant ($p<0.01$) and negative treatment effect, in line with the previous Results 1 and 2. However, the coefficient of the interaction term shows that there is an age-specific positive treatment effect, hence working in the opposite direction. This means that older subjects are less reluctant to also report the prize winning outcome to benefit somebody else, instead of themselves.

Assuming that the prize to win is less attractive for older subjects, this makes intuitive sense: On the one hand, younger subjects, who are keen on winning the prize themselves, are less wiling to report the prize winning outcome to benefit somebody else. On the other hand, older subjects, who are not so keen on winning the prize themselves, may rather see the opportunity to do somebody else a favor. This leads to our next result:

Result 3. *The size and direction of the treatment effect depend on the age of the subject. The treatment effect is largest among very young subjects and decreases in age. Among the oldest subjects, the effect turns positive: The prize winning outcome is more likely reported when the receiver is another child.*

6 Conclusion

In this study we investigated whether children adapt their cheating behavior, dependent on who would benefit from cheating. By changing the recipient of the benefit, we ruled out

that cheating happens out of selfish motives. With our results, we contribute to a growing literature on dishonest behavior, shedding light on motives for cheating that may originate during childhood.

We ran a lab in the field experiment with pairs of children as participants. Participants could benefit themselves through cheating (control group), or benefit the partnering child instead (treatment group). We found that there is relatively more upward compared to downward cheating when the participants can benefit themselves, or when they can benefit somebody who is not a stranger. With this finding we contribute evidence, that while most children are either honest or motivated to cheat out of selfishness, some children do cheat out of other motives: At least when the other party was known to them, a significant number of children was altruistically motivated and cheated in favor of someone else.

Moreover, our results contribute evidence in line with existing literature claiming that very young children are self-centered and only start understanding moral concepts such as fairness or equality as they age. This result is shown by the size of the treatment effect decreasing in age. A few observations from older children, who seemed less attracted by the prize, suggest that some may even have engaged in downward cheating to avoid receiving the benefit themselves. Disentangling shares of upward and downward cheating remains an open challenge and limits the explanatory power of our results. More generally, future research is necessary to closer examine the motives of downward cheating, and under which circumstances such behavior, present among adults, is picked up during childhood.

References

Abeler, J., A. Becker, and A. Falk
2014. Representative evidence on lying costs. *Journal of Public Economics*, 113:96–104.

Andreoni, J., B. Erard, and J. Feinstein
1998. Tax compliance. *Journal of Economic Literature*, 36(2):818–860.

Andreoni, J. and L. Vesterlund
2001. Which is the fair sex? gender differences in altruism. *The Quarterly Journal of Economics*, 116(1):293–312.

Barkan, R., S. Ayal, and D. Ariely
2015. Ethical dissonance, justifications, and moral behavior. *Current Opinion in Psychology*, 6(DEC):157–161.

Belot, M. and J. Van de Ven
2011. Friendships and favouritism on the schoolground–a framed field experiment. *The Economic Journal*, 121(557):1228–1251.

Bucciol, A. and M. Piovesan
2011. Luck or cheating? a field experiment on honesty with children. *Journal of Economic Psychology*, 32(1):73–78.

Buchmann, M., M. Gummerum, M. Keller, and T. Malti
2009. Children's moral motivation, sympathy, and prosocial behavior. *Child Development*, 80(2):442–460.

Chen, Y. and S. X. Li
2009. Group identity and social preferences. *American Economic Review*, 99(1):431–57.

Chytilova, J. and V. Korbel
2014. Individual and group cheating behavior: a field experiment with adolescents. Technical report, IES Working Paper.

Conrads, J., B. Irlenbusch, R. M. Rilke, and G. Walkowitz
2013. Lying and team incentives. *Journal of Economic Psychology*, 34:1–7.

Conrads, J. and S. Lotz
2015. The effect of communication channels on dishonest behavior. *Journal of Behavioral and Experimental Economics*, 58:88–93.

Dreber, A., E. von Essen, and E. Ranehill
2014. Gender and competition in adolescence: task matters. *Experimental Economics*, 17(1):154–172.

Eckel, C. C. and P. J. Grossman
1998. Are women less selfish than men? evidence from dictator experiments. *Economic Journal*, 108(448):726–35.

Erat, S. and U. Gneezy
2012. White lies. *Management Science*, 58(4):723–733.

Fareri, D. S., M. A. Niznikiewicz, V. K. Lee, and M. R. Delgado
2012. Social network modulation of reward-related signals. *Journal of Neuroscience*, 32(26):9045–9052.

Fehr, E., D. Glätzle-Rützler, and M. Sutter
2013. The development of egalitarianism, altruism, spite and parochialism in childhood and adolescence. *European Economic Review*, 64:369–383.

Fischbacher, U. and F. Föllmi-Heusi
2013. Lies in disguise - an experimental study on cheating. *Journal of the European Economic Association*, 11(3):525–547.

Gino, F., S. Ayal, and D. Ariely
2013. Self-serving altruism? the lure of unethical actions that benefit others. *Journal of Economic Behavior & Organization*, 93:285–292.

Glätzle-Rützler, D. and P. Lergetporer
2015. Lying and age: An experimental study. *Journal of Economic Psychology*, 46:12–25.

Gneezy, U.
2005. Deception: The role of consequences. *American Economic Review*, 95(1):384–394.

Houser, D., J. A. List, M. Piovesan, A. Samek, and J. Winter
2016. Dishonesty: From parents to children. *European Economic Review*, 82:242–254.

Houser, D. and D. Schunk
2009. Social environments with competitive pressure: Gender effects in the decisions of german schoolchildren. *Journal of Economic Psychology*, 30(4):634–641.

Houser, D., S. Vetter, and J. Winter
2012. Fairness and cheating. *European Economic Review*, 56(8):1645–1655.

Irlenbusch, B. and M. C. Villeval
2015. Behavioral ethics: how psychology influenced economics and how economics might inform psychology? *Current Opinion in Psychology*, 6:87–92.

Jacob, B. A. and S. D. Levitt
2003. Rotten apples: An investigation of the prevalence and predictors of teacher cheating. *The Quarterly Journal of Economics*, 118(3):843–877.

Jiang, T.
2013. Cheating in mind games: The subtlety of rules matters. *Journal of Economic Behavior and Organization*, 93:328–336.

Killen, M. and J. Smetana
2008. Moral judgment and moral neuroscience: Intersections, definitions, and issues. *Child Development Perspectives*, 2(1):1–6.

Kimbrough, E. O. and J. P. Reiss
2012. Measuring the distribution of spitefulness. *PloS one*, 7(8):e41812.

Klein, S. A., I. Thielmann, B. E. Hilbig, and I. Zettler
2017. Between me and we: The importance of self-profit versus social justifiability for ethical decision making. *Judgment & Decision Making*, 12(6).

Maggian, V. and M. C. Villeval
2016. Social preferences and lying aversion in children. *Experimental Economics*, 19(3):663–685.

Mazar, N., O. Amir, and D. Ariely
2008. The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research*, 45(6):633–644.

McCabe, D. L. and L. K. Trevino
1997. Individual and contextual influences on academic dishonesty: A multicampus investigation. *Research in Higher Education*, 38(3):379–396.

McCabe, D. L., L. K. Trevino, and K. D. Butterfield
2001. Cheating in academic institutions: A decade of research. *Ethics & Behavior*, 11(3):219–232.

Nonis, S. and C. O. Swift
2001. An examination of the relationship between academic dishonesty and workplace dishonesty: A multicampus investigation. *Journal of Education for Business*, 77(2):69–77.

Nunner-Winkler, G. and B. Sodian
1988. Children's understanding of moral emotions. *Child Development*, Pp. 1323–1338.

Nye, J. S.
1967. Corruption and political development: A cost-benefit analysis. *American Political Science Review*, 61(2):417–427.

Pittarello, A., M. Leib, T. Gordon-Hecker, and S. Shalvi
2015. Justifications shape ethical blind spots. *Psychological Science*, 26(6):794–804.

Rosenbaum, S. M., S. Billinger, and N. Stieglitz
2014. Let's be honest: A review of experimental evidence of honesty and truth-telling. *Journal of Economic Psychology*, 45:181–196.

Shalvi, S., F. Gino, R. Barkan, and S. Ayal
2015. Self-serving justifications: Doing wrong and feeling moral. *Current Directions in Psychological Science*, 24(2):125–130.

Stigler, G. J.
1962. Information in the labor market. *Journal of Political Economy*, 70(5, Part 2):94–105.

Thielmann, I. and B. E. Hilbig
2018. No gain without pain: The psychological costs of dishonesty. *Journal of Economic Psychology*.

Tyler, T. R. and S. L. Blader
2005. Can businesses effectively regulate employee conduct? the antecedents of rule following in work settings. *Academy of Management Journal*, 48(6):1143–1158.

Utikal, V. and U. Fischbacher
2013. Disadvantageous lies in individual decisions. *Journal of Economic Behavior & Organization*, 85:108–111.

Whitley, B. E.
1998. Factors associated with cheating among college students: A review. *Research in Higher Education*, 39(3):235–274.

Wiltermuth, S. S.
2011. Cheating more when the spoils are split. *Organizational Behavior and Human Decision Processes*, 115(2):157–168.

Appendix

Figure 4: entertaining a child in front of the experimental area

Figure 5: outside the experimental area

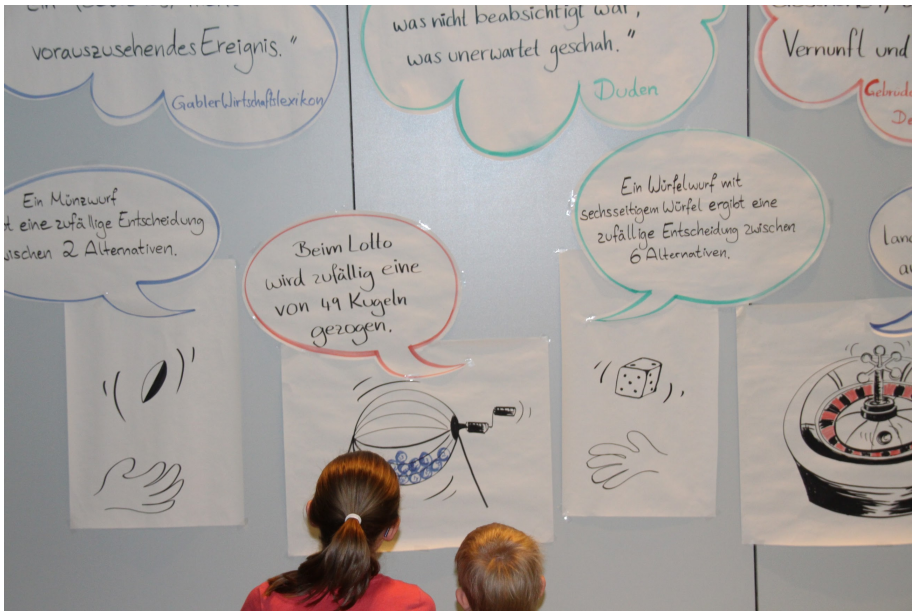

Figure 6: wall with pictures and information on random devices

Figure 7: two-color die to randomly determine roles in the experiment

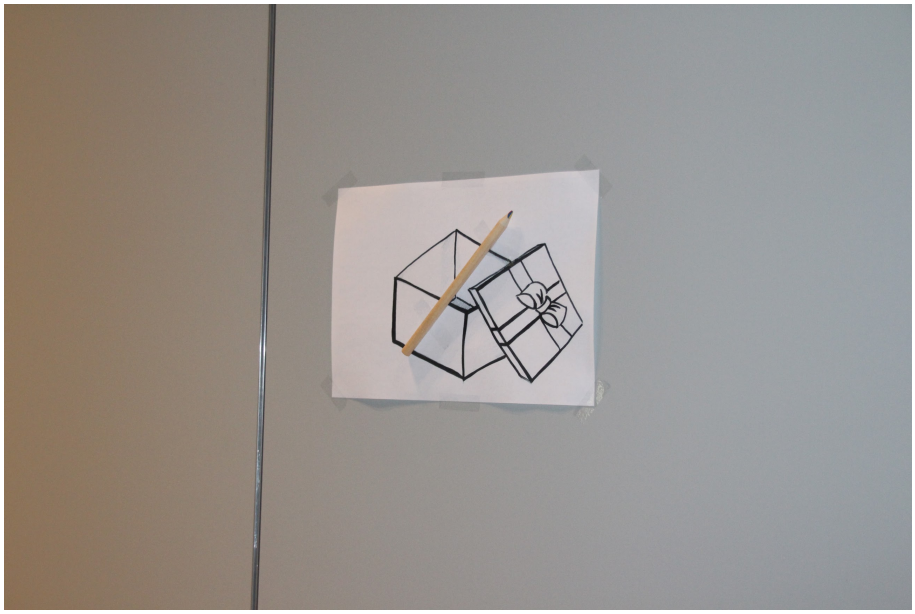

Figure 8: rainbow-colored pencil

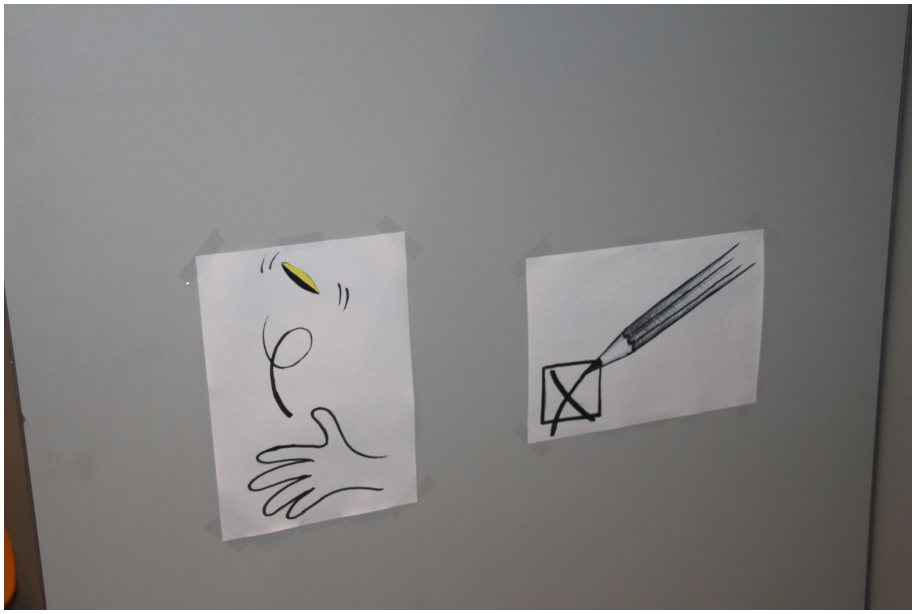

Figure 9: graphical instructions within the private corner 1

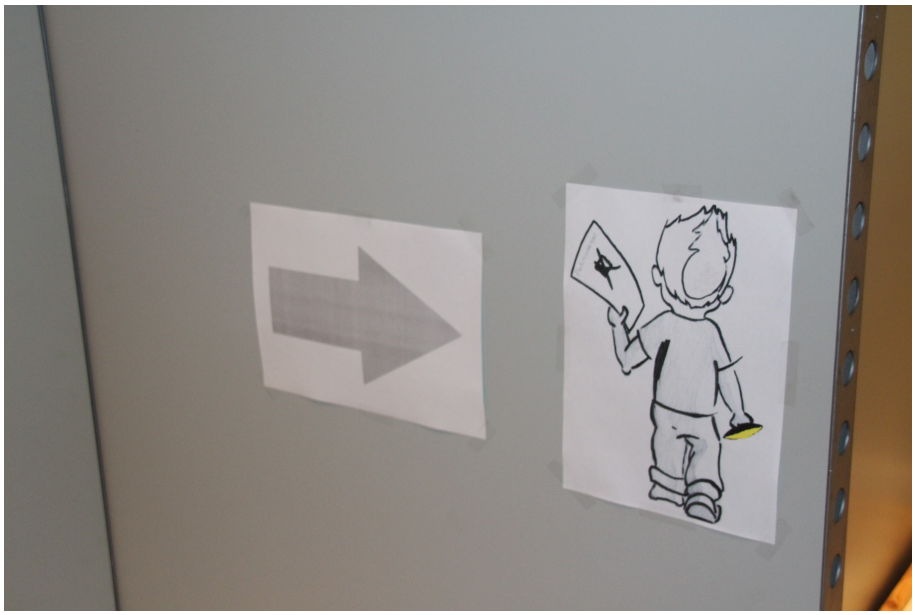

Figure 10: graphical instructions within the private corner 2

	Additional Probit Estimates
treatment	-0.108
	(0.08)
age_subject	-0.012
	(0.01)
female_partner	0.086
	(0.07)
siblings	-0.031
	(0.08)
female_subject	0.112
	(0.09)
female_treat	-0.137
	(0.12)
N	250

Table 3: Probit marginal effects with reporting "yellow" as dependent variable. Control variables include the subject's age, the partner's gender and controls for whether the subject and partner are siblings. Other than in Table 2, here we also include a control variable for the subject's gender, and an interaction between treatment and subject gender (female_treat). * p<0.1, ** p<0.05, *** p<0.01, (se)

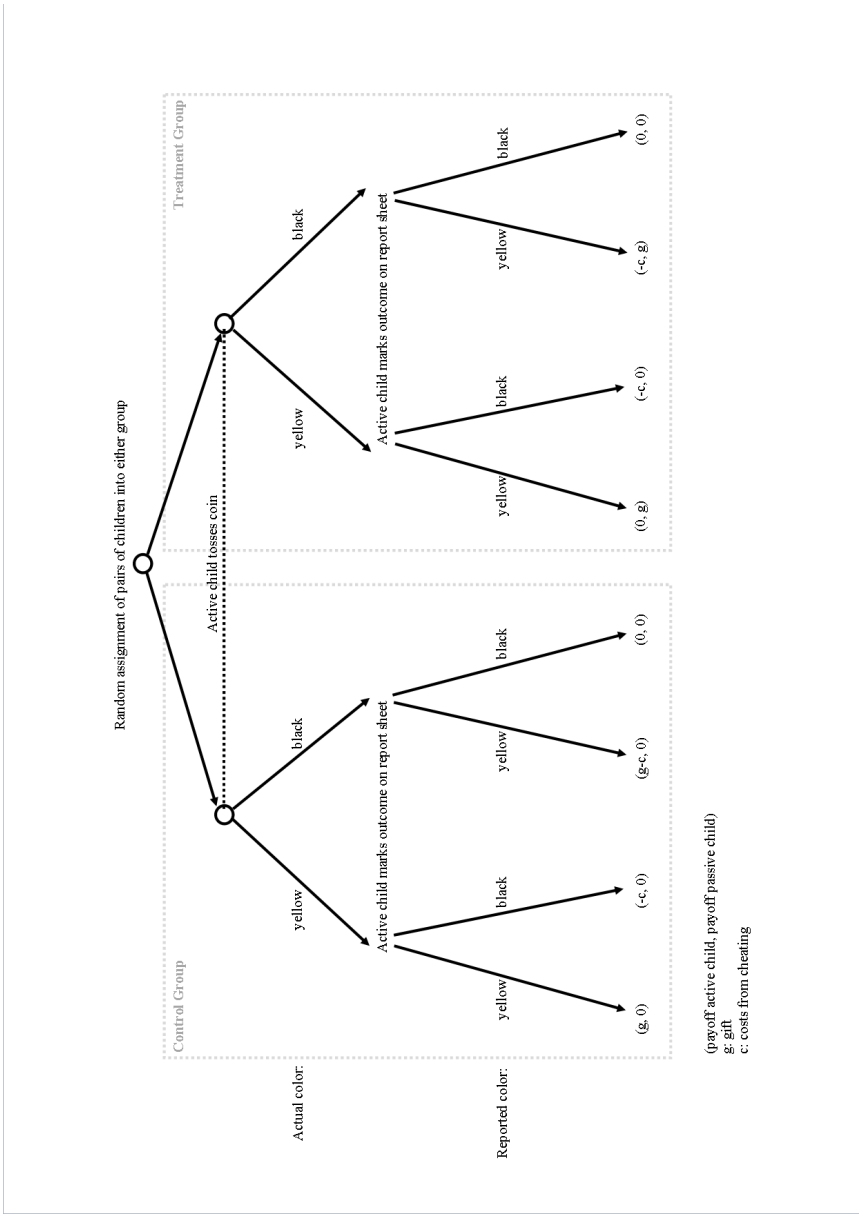

Figure 11: game tree with payouts for participating subjects and passive partners