

On Fraud, Ignorance, and Lying:
Essays in Behavioral Business Ethics

Inauguraldissertation

zur

Erlangung des Doktorgrades

der

Wirtschafts- und Sozialwissenschaftlichen Fakultät

der

Universität zu Köln

2014

vorgelegt von

Diplom-Volkswirt Julian Conrads

aus

Marburg

Referent: Professor Dr. Bernd Irlenbusch
Korreferent: Professor Dr. Matthias Sutter
Tag der Promotion: 23.07.2014

Danksagung

Ich danke meinem Doktorvater Bernd Irlenbusch für die spannenden und erkenntnisreichen Jahre meiner Promotionszeit. Besonders dankbar bin ich ihm für die kollegiale Zusammenarbeit und die inspirierenden Diskussionen, ebenso wie für das entgegengebrachte Vertrauen, eigene Forschungsideen voranzutreiben und umzusetzen zu können.

Matthias Sutter möchte ich herzlich für hilfreiche Kommentare und die Übernahme des Korreferats danken.

Ganz besonderer Dank gebührt meinen Lehrstuhlkollegen Rainer Michael Rilke und Gari Walkowitz. Der Aufbau des Lehrstuhls, die gemeinsamen Forschungsprojekte und die zahlreichen wissenschaftlichen und nicht-wissenschaftlichen Gespräche haben meine Promotionszeit zu einer wunderbaren Erfahrung gemacht. Ebenso danke ich Felix Ebeling und Sebastian Lotz für die gelungene Zusammenarbeit, die zur Studie des fünften Kapitels dieser Dissertation geführt hat.

Sehr dankbar bin ich auch allen weiteren Ko-Autoren, Weggefährten und Unterstützern, die bei der Entstehung dieser Arbeit direkt oder indirekt beteiligt waren: Johannes Berger, Anja Bodenschatz, Anastasia Danilov, Patrick Kampkötter, Andrew Kinder, Lisa Klinger, Felix Kölle, John List, Caro Martens, Janna Ter Meer, Frauke Meyer, Mattia Nardotto, Albenä Neschen, Axel Ockenfels, Katharina Peeters, Katrin Recktenwald, Tommaso Reggiani, Marina Schröder, Anne Schielke, Dirk Sliwka, Patrizia Stumper, Philipp Tillmann, Dagmar Weiler, Peter Werner und Daniel Wiesen.

Meinen Freunden außerhalb der akademischen Welt bin ich für die erholsamen Ablenkungen und die ‘anderen Blickwinkel’ sehr verbunden, sei es beim Sport, Musizieren oder sonstigen Freizeitaktivitäten.

Herzlicher Dank gilt auch meiner Familie, die mich während des Studiums und der Promotionszeit immer gefördert und in meinen Vorhaben bestärkt hat.

Unbeschreiblich dankbar bin ich schließlich meiner Lebenspartnerin Babett Rutsch, die mich bedingungslos und liebevoll unterstützt hat, und dabei ein wunderbares Gespür dafür hatte, mir zum richtigen Zeitpunkt die wirklich wichtigen Dinge des Lebens ins Gedächtnis zu rufen.

Contents

1	Introduction	1
2	Strategic Ignorance in Ultimatum Bargaining	8
2.1	Introduction	8
2.2	Related Literature	11
2.3	Hypotheses	13
2.4	Experimental Design and Procedures	16
2.4.1	A Simple Framework of Strategic Ignorance in Bargaining	16
2.4.2	Treatments	18
2.4.3	Procedures	19
2.5	Results	21
2.5.1	Results <i>Transparency</i> and <i>Non-Transparency</i>	21
2.5.2	Results <i>Choice</i> Treatments	23
2.5.3	Proposer Payoffs	26
2.5.4	Hypothetical Decisions and Beliefs in <i>Choice</i>	28
2.6	Concluding Remarks	29
2.7	Appendix	32
3	Lying and Team Incentives	40
3.1	Introduction	40
3.2	Experiment	42
3.3	Hypotheses	45
3.4	Results	45
3.5	Concluding Remarks	52
3.6	Appendix	53
4	The Effect of Communication Channels on Lying	55
4.1	Introduction	55

4.2	Experiment	59
4.2.1	Experimental Design	59
4.2.2	Experimental Procedures	60
4.3	Theoretical Considerations and Behavioral Hypotheses	61
4.4	Results	64
4.4.1	Behavior of Subjects	64
4.4.2	Beliefs about Behavior of Other Subjects	67
4.5	Discussion: Partial Lying and Communication Channels	69
4.6	Conclusion	70
4.7	Appendix	72
5	Fraud at the Kiosk - Field Evidence from Buying Candy	78
5.1	Introduction	78
5.2	A Brief Description of Kiosks	80
5.3	Experiment 1: Overcharging When Selling Candy	82
5.3.1	Procedure	82
5.3.2	Results and Discussion	83
5.4	Experiment 2: Behavioral Consistency of Sellers and Discrimination of Buyers	85
5.4.1	Procedure	86
5.4.2	Results and Discussion	86
5.5	General Discussion	90
5.6	Appendix	93
	Bibliography	97

List of Tables

2.1	Results from <i>Transparency</i> and <i>Non-Transparency</i>	22
2.2	Expected payoffs (<i>EP</i>) and observed total surplus	27
A2.1	Proposers' offers and receivers' acceptance rates	34
A2.2	OLS and probit regressions on decisions in state s_n	35
3.1	Descriptive statistics	47
3.2	Explaining reported production output p_i	50
4.1	Treatment variations	60
4.2	Overview of results	65
4.3	Regression analysis of reported outcomes	67
A4.1	Probit regressions on reporting the outcome 4	72
A4.2	Additional regression analysis of reported outcomes	73
5.1	Overview of results from Experiment 1	84
5.2	Treatments of Experiment 2	86
5.3	Overview of results from Experiments 1 and 2	87
5.4	OLS regressions on Δ_i for Experiments 1 and 2	88
5.5	Classification of predominant seller types	89
A5.1	Overview of sellers' demographics and personality traits	93
A5.2	Amount of overcharging in cents	94
A5.3	Amount of undercharging in cents	94
A5.4	OLS regressions for Experiment 1	95
A5.5	Weight estimation of candy bags in laboratory	95
A5.6	Results from Experiments 1 and 2	96

List of Figures

2.1	Game trees in <i>Transparency</i> and <i>Non-Transparency</i> . . .	17
2.2	Percentages of proposers offering option A in state s_n . .	24
2.3	Percentages of responders accepting option A in state s_n	25
A2.1	Game-tree of the <i>Choice</i> treatment	32
A2.2	Game-tree of the <i>Choice Uncertain Information Acquisition (Choice-UI)</i> treatment	32
A2.3	Game-tree of the <i>Choice Hidden (Choice-H)</i> treatment .	33
3.1	Relative frequency of reported production outputs in the three different treatments	46
4.1	Frequencies of reported tails	64
4.2	Actual behavior and beliefs on payoff-maximizing reports	68
A4.1	Frequencies of reported tails by men	74
A4.2	Frequencies of reported tails by women	74
A4.3	Beliefs about frequencies of reported tails	75
A5.1	Map of kiosks	93

1 Introduction

This thesis empirically investigates deceptive behavior in social interactions. It includes four experimental studies dealing with fraud, ignorance, and lying. The aim is to better understand what drives these types of deceptive behavior and how they are influenced by the design of the decision-making environment. From a theoretical perspective, most of standard economic models assume that people deceive if it is to their own benefit. People thereby weigh the benefits from deception against the potential cost of being detected and punished (e.g., Becker, 1968, on crime). More recent models, however, assume that some people face a sort of intrinsic costs that prevent them from engaging in deceptive acts (Kartik, 2009, and Gibson et al., 2013). From a practical perspective, the need to better understand mechanisms that lead to deceptive conduct is not only evident from various high-profile scandals in corporations and public institutions, but also from the prevalence of dishonest behavior in everyday life. Although media coverage typically focuses on sensational frauds (e.g., those run by Deutsche Bank or Bernard Madoff), daily “ordinary dishonesty” is nevertheless widespread, for example, misrepresentation about one’s own performance at work or manipulating one’s annual tax declaration. One robust empirical finding reveals that not only are a few bad apples rotten, but rather that many apples in the barrel are just a little bit bad (Bazerman and Tenbrunsel, 2011). Taken cumulatively, the damages from dishonest behavior for corporations and society as a whole are considerable (Feldman, 2009). An overarching goal of this thesis is to provide empirical insights into potentially adverse effects of established frameworks, such as corporate structures or incentives schemes, and to propose possible solutions for organizational designers and businesses.

This dissertation focuses on self-serving deception that results in personal profits. Well-intentioned deception (e.g., about a friend’s new haircut), which is also an important part of social interactions, will not be

analyzed. In the four studies presented, two basic requirements for the possibility of deception are at hand: asymmetric information between two parties and communication among parties (Croson, 2005). The first study concerns (strategic) ignorance in a bargaining environment where the proposing party has the opportunity to deliberately remain ignorant about the position of another party in order to selfishly maximize personal profits. In this case, deception arises as a result of the proposer's conscious decision not to acquire information. This omission is made with the hope that even unfavorable offers will be accepted by the other party. The second and third studies employ two different perspectives to examine how personal profits can be maximized through dishonesty, but focus on direct lying about private information. In the second study, the prevalence of lying in relation to different performance-related incentive schemes is examined, and in the third study the effect of different channels of communication on the individual propensity to lie is analyzed. Finally, the fourth study addresses fraudulent behavior of sellers in a real-world market, in which the sellers have an informational advantage over their customers.

The scientific approach applied in this dissertation follows that of behavioral ethics, which is concerned with explaining individual behavior that occurs in the context of larger social prescriptions (see Tenbrunsel and Smith-Crowe, 2008). In contrast to the normative philosophical approach, which focuses on the question of how people should act, behavioral ethics is a descriptive facet of ethics and might be better suited for understanding the root causes of unethical behavior (see Bazerman and Gino, 2012). From a methodological perspective, behavioral ethics applies experimental techniques originating from social psychology and experimental economics. While it is particularly difficult to observe unethical conduct and its influencing factors in the field, e.g., people usually try to disguise their unethical misdeeds, an experimental technique allows actions to be perfectly monitored. The influence of the decision-making environment and contextual differences can be systematically varied to better understand observed behavior. The most important advantage of the experimental approach is control, which is essential when aiming for causal interpretations of (unethical) behavior (see Falk and Heckman, 2009). In philosophy, experimental investigations of actual behavior have

been rather neglected (see Knobe and Nichols, 2008, and Greene, 2012, for exceptions).

The following paragraphs briefly outline the separate research questions of each chapter and summarize the main findings and contributions.

Chapter 2 examines strategic ignorance in a bargaining context.¹ It addresses the research question of whether it pays off to remain strategically ignorant about the position of another bargaining party to achieve the highest possible payoff for oneself. The idea of introducing the option of remaining deliberately ignorant about the position of another party - which can either be wealthy or poor - is adapted from Dana et al. (2007). The authors analyze the strategic use of ignorance in dictator-game giving and show that many dictators take advantage of this strategic “moral wiggle room” to justify a selfish actions. When transferring this approach to a bargaining context, the proposing party has the option to remain ignorant about the position of the other party, with the hope that the latter will accept even an unfavorable offer. The rationale is that the accepting party is more likely to agree to an offer proposed under ignorance. This strategy would contradict the general assumption that more information in bargaining is always better (e.g., Fisher and Ury, 1991). In his classic article “An Essay on Bargaining”, Schelling (1956) introduced the argument that ignorance might actually be a strength rather than a weakness. We test and confirm this conjecture in a simple take-it-or-leave-it ultimatum bargaining experiment, in which a proposer can choose between two possible options to offer. One option always gives the proposer a higher payoff than the other. The payoff for the responder depends on the (randomly determined) state of nature. In one state payoffs are aligned and in the other they are not. Thus, an ignorant proposer who offers his preferred option does not know at the outset whether this option is also favorable to the responder. In one treatment variation we find that offers - even unfavorable ones - by a proposer who is by default ignorant about the state of nature are almost always accepted by the responder. In additional treatments, we introduce the possibility of remaining intentionally ignorant. About one quarter of proposers choose to be ignorant. Unfavorable offers from a proposer who deliberately chooses to be ignorant are more frequently accepted than comparable offers by a proposer who

¹Chapter 2 is based upon Conrads and Irlenbusch (2013).

informed himself. A critical feature of unfavorable offers proposed under ignorance is the proposer’s intentionality. Responders seem to accept unfavorable offers proposed under ignorance because negative intentions are less present compared to situations in which unfavorable offers are knowingly proposed (see Blount, 1995; Falk et al., 2008). In summary, we present evidence that a proposer’s informational weakness in bargaining can be a strength, as hypothesized by Schelling (1956). In addition, we can extend his conjecture by showing that ignorance can even be used strategically if the other bargaining party is aware of the proposer’s ignorance but uncertain whether the ignorance was intentionally chosen.

Chapters 3 and 4 of this thesis investigate lying behavior from an organizational perspective. Lying has become a lively area of investigation in the experimental economics literature (see Croson, 2005). It has been demonstrated that a large share of people misreport private information to their own material advantage. Gneezy (2005), for example, shows that people lie more often as the personal benefit from lying increases and the loss for the deceived person decreases. In many situations, however, people can decide whether to lie to some or the full extent. Fischbacher and Föllmi-Heusi (2013) show that a high share of people lie ‘incompletely’ in order to disguise the lie and appear honest to themselves and others (see also Mazar et al., 2008). Theoretical models try to capture these behavioral patterns by considering heterogeneous (psychosocial) moral costs of lying (see Kartik, 2009; Gibson et al., 2013).

A growing strand of research has demonstrated the existence of unethical conduct under different types of incentive schemes, for example, under tournament incentives (e.g., Harbring and Irlenbusch, 2011; Conrads et al., 2014) or goal-setting (e.g., Schweitzer et al., 2004). Chapter 3 compares the influence of two widely-used compensation schemes, individual piece rates and team incentives, on individuals’ inclination to lie.² Both incentive schemes evidently possess effort-enhancing effects (Lazear and Gibbs, 2009). Under team incentives, revenues from producing an output are shared among agents who build a team, whereas under individual piece rates revenues are not shared. The manner in which these two incentive schemes induce people to engage in lying behavior is a rather neglected feature of the existing literature. There

²Chapter 3 is based upon Conrads et al. (2013).

are at least two possible explanations for the higher incidence of cheating under a team incentive scheme compared to individual piece rates. First, gains from lying can be split among team members, which offers an easy justification to lie. Second, lying is harder to detect since individual responsibility is diffused, i.e., in most cases deceptive acts cannot be unambiguously attributed to individuals (see Bandura et al., 1996). Adapting the experimental die-rolling setup of Fischbacher and Föllmi-Heusi (2013), we analyze different treatments in which individual and team incentives are employed. Lying turns out to be more pronounced under team incentives than under individual piece rates. When disentangling the different explanations of lying under team incentives, we find that subjects tend to lie more under team incentives because they can diffuse their responsibility for lying. Therefore, from a practical perspective, our study provides designers of incentive systems empirical evidence into the potentially harmful side effects of a presumably effort-enhancing team compensation scheme.

Chapter 4 focuses on piece rate compensation and investigates the effect of different channels of communication on lying. From a practical perspective, this question is of interest since communication has shifted from personal direct communication, e.g., face-to-face, to more anonymous indirect computer-mediated communication, e.g., online. This change in communication context may also affect norms of social interactions. Increasing social and physical distance of an interaction may weaken the social norm not to lie (see Bohnet and Frey, 1999; Charness and Gneezy, 2008). Within experimental economics literature, the influence of different communication media is analyzed in social dilemma games (e.g., Brosig et al., 2003) and bargaining games (e.g., Valley et al., 1998). Both studies find higher degrees of cooperation and trust under face-to-face pre-play communication. The study in chapter 4 systematically varies the channel of communication to analyze its pure effects on lying. By adapting an experimental coin-flip game, private information about a random production output is either communicated face-to-face, by phone, via computer, or online. In the experimental setup, subjects have the chance to lie either to some or the full extent to increase their personal payoffs. The results suggest that the majority lies independent of the means of communication. However, the decision whether to lie

fully or partially strongly depends on the communication channel employed. The smaller the social distance of the communication medium is, the less subjects lie maximally. This study therefore contributes to other findings on incomplete lying and adds empirical evidence on the effect of different communication channels. From a practical view, it is important to consider which channels of communication are best suited to specific social and business interactions, e.g., in e-business, as the use of some channels may backfire in certain contexts.

In the fifth chapter of this dissertation, we step out of the laboratory and investigate fraudulent behavior in the field (see List, 2006, for an overview on the general advantages of field experiments).³ The study addresses whether and how sellers exploit their informational asymmetries vis-à-vis customers by defrauding them. This research is related to the literature on credence goods markets, in which customers cannot judge - either ex ante or ex post - whether the type or quality of the good is the one they actually needed (see Darby and Karni, 1973; Dulleck et al., 2011). Credence goods have been most prominently discussed in the realms of auto repair or medical services, where customers usually suffer from inferior knowledge compared to expert sellers. Typically, expert sellers in these markets can either overcharge or overtreat their customers. The market investigated in this chapter is not a typical credence goods market but shares its essential characteristics, i.e., sellers have an informational advantage over their customers that allows for fraudulent overcharging, but not overtreatment. More precisely, we analyze behavior of sellers in urban kiosks who sell loose candy according to candy-by-weight pricing. Some sellers secretly weigh the candy behind the sales counter, while others weigh on a scale publicly visible to the customer. In our first experiment, we find that weighing secretly leads to significant overcharging compared to when weighing is observable by the customer. In our second experiment, we implement a panel in the field by conducting four consecutive purchases in kiosks with hidden scales to analyze the behavioral consistency of sellers. In addition, we manipulate the appearance of the buyer by making him appear either rich or poor, and also varying the amount of candy bought, i.e., either a small or a large bag of candy (the experimental design is inspired by Balafoutas

³Chapter 5 is based upon joint work with Felix Ebeling and Sebastian Lotz.

et al., 2013). Our findings lead to the conclusion that sellers display a behavioral consistency that enables categorization of sellers into accurate or fraudulent. In addition, we show that sellers tend to discriminate against high status buyers who purchase large quantities. The experiments provide initial results on the consistency of overcharging behavior and can therefore inform recent hypotheses surrounding credence goods markets. As a practical solution against overcharging, we can show in this context that transparency of prices completely mitigates the risk for customers, as overcharging is negligible when scales are publicly visible.

In summary, the four experimental studies of this dissertation reveal that a considerable share of people engage in deceptive behavior. However, honesty and incomplete lying can also be observed, which may underline the existence of heterogeneous preferences for truthfulness. The following four chapters present the separate research studies in detail.

2 Strategic Ignorance in Ultimatum Bargaining

2.1 Introduction

The availability of information on an opponent's bargaining position plays an important role in negotiations and not only affects one's own bargaining behavior but also the behavior of an opponent. Generally, it is assumed that the more information that is available in a bargaining situation, the better the bargaining position is (e.g., Fisher and Ury, 1991, p.45). Schelling (1960) challenged this view by arguing that a bargainer who is incompletely informed about his opponent's payoffs might have an advantage because the opponent would be forced to make concessions to avoid a bargaining breakdown. In his chapter on "*Strategic Moves*", Schelling notes, "*(...) ignorance can be an advantage to a player if it is recognized and taken into account by an opponent*" (Schelling, 1960, p.161). As the informed bargainer is aware that the uninformed one does not know what a reasonable solution is, the burden of avoiding a stalemate is on the side of the informed bargainer. Early experimental studies seem to support this view (Siegel and Fouraker, 1960; Hamner and Harnett, 1975). The following example illustrates the basic intuition: two persons walking on a crowded main street are going to collide. One person anticipates the upcoming event but the other does not, for example, due to a distraction. The person aware of the possible collision clears the way, accepting the "cost" of leaving his ideal route. The other (unintentionally) ignorant person continues along his intended way: being uninformed pays off. Ignorance might also be used strategically. A person who anticipates the possibility of a collision might simply walk down the street while looking at the ground and pretending to be ignorant. The other informed person has to bear the costs of avoiding the

collision, although he might have the feeling that the ignorant person is intentionally avoiding looking up. Thus, remaining strategically ignorant might also pay off. Putting this in an organizational context, one might consider a business partnership. One day an urgent request comes in, but only one of the two partners is in the office. Sub-tasks have to be allocated quickly between the two partners, and the nature of the tasks prohibits subsequent re-allocation. By deliberately remaining ignorant and not asking her partner about his preferences, the partner in the office can pick her preferred sub-tasks and leave the other sub-tasks to her partner. Should the partner turn out to dislike the sub-tasks allocated to him, she can come up with the excuse: *“Oh sorry, I didn’t know”*. The excuse might still have some force despite the fact that, in principle, she could have informed herself – or at least attempted to do so – for example, by giving her partner a call.¹

The aim of this study is to experimentally test Schelling’s conjecture in a simple two-person take-it-or-leave-it bargaining game. As it is particularly difficult to observe (strategic) ignorance in bargaining in the field, we chose an experimental approach that allows actions to be perfectly monitored, including those in which one attempts to avoid acquiring information. Control is the most important advantage of an experimental study (see Roth, 1995; Falk and Fehr, 2003), which is essential for our purpose, i.e., drawing conclusions regarding how strategic ignorance causally affects behavior. Moreover, in contrast to questionnaire studies, it is possible to provide participants with incentives that are likely to have a crucial influence on strategic ignorance in bargaining. Our basic experimental framework comprises a simple situation that is reduced to the essential features of strategic ignorance. One of two states of nature is determined by a 50:50 draw. While the interests of a proposer and a responder are aligned in state s_a , they are not in state s_n . The proposer has to offer one of two options, option A or option B . In state s_n , the proposer profits from option A more than the responder. Option B in

¹Fischbacher and Utikal (2013) analyze the effectiveness of apologies in preventing punishments after harmful offenses. They find that excuses are not accepted if the harm-doer commits offenses intentionally. If the intention of an offense is not clear, i.e., if the situation is ambiguous, apologies seem to be an effective instrument to reduce the likelihood of being punished. In our context, remaining ignorant blurs the intentionality of the proposer and therefore might reduce the likelihood of being punished with a rejected offer.

state s_n would make both players' payoffs nearly equal, but this option is slightly inferior for the proposer in comparison to option A . In state s_a , option A provides both players with higher payoffs than option B . The responder can accept or reject the offer. Accepting an offer always leads to positive payoffs for both players, while rejection leaves them with zero payoffs.

In treatment *Transparency*, both players are fully informed about the true state, and we observe that proposers are not always able to implement their most preferred option. Offers of option A are frequently rejected in state s_n . In the *Non-Transparency* treatment, the proposer is ignorant about the true state, but the responder knows it. This information is known to both players. We hypothesize that the proposer will benefit from being ignorant, as the responder will accept nearly all offers. As the experimental results show, an ignorant proposer can almost always implement her most preferred option, i.e., option A . A possible explanation for this result is differences in causal attributions of how the outcomes emerged. If an unfavorable offer is attributable to bad luck (i.e., the random choice of one of the two states of nature), responders might accept these offers because negative intentions are not involved (Blount, 1995; Falk et al., 2008).

In a third treatment, *Choice*, the proposer can choose between remaining ignorant about the state of nature or inform herself about it. None of the alternatives incur any direct monetary costs. The notion of introducing the possibility of remaining strategically ignorant of the opponent's payoff is adapted from Dana et al. (2007), who analyze the strategic use of ignorance in a dictator game. The dictator can remain ignorant to justify a selfish action to herself. In our setting, not to inform herself about the state also allows the proposer to select the self-interested offer (i.e., option A) without knowing the actual payoff consequences for the responder. Knowing the state would potentially place some (internal) pressure on the proposer to select the more equalizing option B in state s_n . Additionally, by remaining ignorant, the proposer might wish to influence the responder's inclination to accept option A in state s_n . The responder is always informed about the actual state and learns whether the proposer chose to remain ignorant. We hypothesize that proposers will not benefit from strategic ignorance, as responders will perceive the act of remaining

ignorant as hostile. Our results suggest that responders tend to reject option A in state s_n less frequently when the proposers remain ignorant. To push the notion of the perception of hostile intentions a bit further, we designed a modified version of the *Choice* treatment, *Choice Uncertain Information Acquisition*, where a proposer’s attempt to inform herself about the state is only successful in 50% of the cases. As a consequence, if the proposer remains ignorant, the responder does not know whether this ignorance was purposeful. We find that responders accept option A offers from ignorant proposers significantly more frequently in state s_n than from proposers who successfully informed themselves about the state. In a fifth treatment, *Choice Hidden*, the responder is not informed of whether the proposer informed herself about the state. Here, few proposers remain ignorant, and responders frequently accept option A offers in state s_n .

The paper is organized as follows. We begin by discussing the literature related to strategic ignorance. Second, we state our hypotheses and elaborate our experimental design. In section five, we report the experimental results. Finally, section six discusses the results in light of previous findings and concludes.

2.2 Related Literature

Proctor and Schiebinger (2008, p.3) emphasize the omnipresence of ignorance and differentiate – from an epistemic perspective – between “*ignorance as native state (or resource), ignorance as lost realm (or selective choice), and ignorance as a deliberately engineered and strategic play (or active construct)*”. In our study, we particularly focus on the third category of ignorance. Although they assess the considerable relevance of strategic ignorance in human interactions, the literature on strategic ignorance in bargaining is relatively small. Some experimental studies have indicated that negotiators might not profit from being uninformed. For example, Roth and Murnighan (1982) showed that varying information asymmetries between negotiators has an impact on how a pie is divided, i.e., uninformed negotiators tend to be exploited by their informed opponents. Negotiators made lower offers if they knew that their opponent

was unaware of the actual size of the pie. Being ignorant turned out to be a disadvantage (see also Kagel et al., 1996).

Other experimental investigations have shown that ignorance might be an advantage. Siegel and Fouraker (1960) conducted a seminal study on the role of ignorance in bargaining. In their bilateral bargaining experiment, the buyer knew the payoff tables of both sides, but the seller only knew his own payoff table. The buyer and seller then had to come up with a price-quantity agreement. Although their results were not statistically significant, the authors found that the incompletely informed participant tended to be better off than the informed opponent. Siegel and Fouraker (1960) argued that the incompletely informed bargainer established a higher aspiration level, as he was unable to form realistic expectations and therefore made larger demands, smaller concessions and accepted longer durations to reach an agreement. A follow up study by Hamner and Harnett (1975) showed a similar effect. Beisecker et al. (1989) examined a complete-incomplete information situation with a fictitious bargaining task. Their results show that an uninformed bargainer can profit from ignorance when the counterpart perceives the own advantage as a violation of procedural equity. To restore relational equity, the completely informed bargainer may accept less favorable agreements. Overall, this strand of literature suggests that ignorance in bargaining can be an advantage. None of these studies, however, examines the possibility of strategically electing to remain ignorant.

More recently, Poulsen and Roos (2010) examined the effect of strategic information avoidance in a Nash demand game where two players had to negotiate over the distribution of a sum of money. At the beginning, the responder had to decide whether he wished to learn about a demand made by a proposer. The proposer was informed about the responder's decision before stating his demand. In a repeated setting, responders learned that more information may be harmful, i.e., over time, information-avoidance increased and the distribution of the surplus became more balanced. In an ultimatum game setup, Poulsen and Tan (2007) allowed the responder to choose a Minimum Acceptable Offer (MAO). The proposer could then costlessly acquire the information about the responder's MAO before making his proposal. The offer was accepted or rejected according to the previously stated MAO. One third of the proposers remained un-

informed and offered half of the pie. Information-acquiring proposers made offers equal to the responders' MAOs. In a treatment without information-acquisition, the MAOs stated by the responders were smaller than those in the information-acquisition treatment, demonstrating that the opportunity to gather information about the MAOs may backfire for the informed party. Thus, in these two studies, one player had the opportunity to remain ignorant of the other player's strategic choice. In our study, proposers are able to remain ignorant of the payoffs resulting from the offer. In a similar vein, Gehrig et al. (2003, 2006) examined a situation in which a proposer could purchase information about a responder's outside option in an ultimatum game. Under transparent information acquisition, where the responder knew whether the proposer was informed, acceptance rates were higher than in non-transparent situations.

As mentioned above, our experimental design is also influenced by the dictator game setup of Dana et al. (2007). The aim of their study was to analyze whether generosity in dictator game giving is truly evidence of a concern for desirable social outcomes. They showed that when the dictator had the opportunity to not know whether his action hurt the receiver, many dictators chose a "moral wiggle room" and made self-interested choices. Dana et al. (2007) concluded that dictators were more concerned with seeming fair to themselves than actually being fair (see also Dana et al., 2006; Broberg et al., 2007).

2.3 Hypotheses

Assuming that it is common knowledge that both bargainers are purely interested in maximizing their own payoffs, a responder should accept any offer that gives him positive payoffs. Thus, a proposer could safely offer a tiny positive amount to the responder while leaving the larger share for herself. However, as we know from numerous experiments on the ultimatum game and other bargaining games, proposers offer more than the smallest positive amount (see, for example, Güth et al., 1982; Camerer, 2003). This seems to be partly driven by the proposer's concerns for (distributive) fairness.² Partly, this behavior results from proposers an-

²Models of social preferences can explain such behavior to some extent. For prominent models of inequity-aversion, see Levine (1998), Fehr and Schmidt (1999) and

anticipating that responders are willing to reject unfair offers and/or offers that signal bad intentions of the proposer (see Blount, 1995; Falk et al., 2008).

As mentioned in the introduction (and will be explained in detail in the next section), two states of nature are possible in our experimental game: a state s_n where interests of a proposer and a responder are in conflict, and a state s_a where interests are aligned. The actual state of nature is randomly determined. The proposer has to offer one of two options to the responder, option A or option B . Option A gives greater payoffs to the proposer than option B , independent of the actual state. Thus, a self-regarding proposer should always offer option A . In state s_a , option A leads to a higher payoff for the responder compared to option B . In state s_n , the opposite is the case. Additionally, in state s_a , option A is the more efficient (in terms of total payoffs) and payoff equalizing option, and in state s_n , option B is the more efficient and payoff equalizing option. The responder has to decide whether to accept or reject the offer. If he rejects the offer, both players receive zero payoffs.

In our experimental design, we vary what a proposer knows about the state of nature. In one setting, she is informed about the actual state. In another, the proposer is kept ignorant of this. In additional settings, the proposer can deliberately decide whether she wants to inform herself about the state, i.e., she can also remain ignorant. The responder always knows the actual state and whether the proposer is informed about the state.³

In our hypotheses, we concentrate on behavior in state s_n , which is the more interesting state for the purposes of our research. In state s_a , payoffs are aligned and it can generally be assumed that option A will be proposed and accepted.

In three treatments of our experimental setup, the proposer can decide whether she wants to inform herself about the state. When unaware of the state, the proposer naturally chooses option A because this gives her a

Bolton and Ockenfels (2000). The influence of intentions is modeled in Falk et al. (2008). Concerns for efficiency might also play a role in our setting (Charness and Rabin, 2002). See also the findings on mini-ultimatum games that employ a reduced strategy set – often two strategies – for the proposer (see, for example, Bolton and Zwick, 1995; Güth et al., 2001; Falk et al., 2003).

³Except in treatment *Choice Hidden*, where the responder does not know whether the proposer informed herself about the state (see details below).

higher payoff, it maximizes expected total surplus, and she does not know the payoff consequences for the responder. If the proposer, however, has efficiency concerns, it would be desirable to know the actual state because option A is only efficient in state s_a and option B is efficient in state s_n . Additionally, if the actual state is s_n , option A could be perceived as being unfair, and the proposer therefore runs the risk that the offer might be rejected. Deliberately not informing oneself might also be perceived as showing of bad intentions. However, based on the results of Dana et al. (2007), some proposers might prefer to exploit the moral wiggle room by remaining ignorant and proposing option A without knowing whether this has adverse consequences for the responder. Additionally, she might believe that it is easier for the responder to accept option A if he knows that the proposer was ignorant of the consequences of her choice for the responder's payoff. Thus, in line with previous experimental findings, the following hypothesis can be derived:

Hypothesis “Information Acquisition of Proposers”: *The majority of proposers inform themselves about the actual state.*

Our second hypothesis concerns the proposer's offer. Let us assume that the proposer knows the actual state. Based on previous findings, we hypothesize that she might be concerned about efficiency and/or equity and therefore chooses option B in state s_n . Additionally, in state s_n , the proposer might fear that the responder will reject option A because this is the less equitable option. This consideration leads to our second hypothesis:

Hypothesis “Proposers’ Offers”: *Proposers who know the actual state tend to offer option B in state s_n . Proposers who do not know the actual state offer option A.*

Our third hypothesis focuses on the responder's acceptance decision. A responder who is primarily concerned with efficiency is not very likely to reject any offer because a rejection would reduce efficiency. Based on previous findings on inequity aversion, we hypothesize that the responder might prefer zero payoffs for both players to the considerably unequal payoff allocation that would result from accepting option A in state s_n . This is true, at least if the proposer knows the actual state. If

the proposer remained ignorant, it is unclear how the responder would react to an offer of option A in state s_n . On the one hand, the responder might be upset that the proposer did not inform herself and therefore reject the offer. On the other hand, the responder might acknowledge that the proposer was ignorant of the consequences of offering option A and therefore accept the offer (this is particularly true if the proposer might have unsuccessfully attempted to inform herself, as could be the case in our treatment *Choice-UI*; see below). If the proposer was unable to inform herself about the true state of nature, ignorance might serve as an excuse, and we therefore hypothesize that in this case the responder is inclined to accept option A in state s_n . Recall that the responder always knows the actual state of nature. Thus, we arrive at the following hypothesis:


***Hypothesis “Responders’ Acceptance Decision”:** Option A offers by proposers who know the actual state tend to be rejected in state s_n . Responders’ reactions to option A in state s_n offered by deliberately ignorant proposers are ambiguous. Option A offers in state s_n by proposers who were unable to acquire information about the actual state tend to be accepted.*

2.4 Experimental Design and Procedures

2.4.1 A Simple Framework of Strategic Ignorance in Bargaining

We consider a non-constant-sum bargaining situation in which a random move selects one of two possible states of nature that occur with equal probability. This is known to all players. A proposer (P) makes a take-it-or-leave-it offer to a responder (R) by choosing between two possible options, option A and option B . The responder has to decide whether to accept or reject the offer. Rejection leaves both players with zero payoffs. Accepted options provide both players with strictly positive payoffs. We focus on take-it-or-leave-it bargaining to keep the interaction simple. A take-it-or-leave-it bargaining structure places the proposer in a position of relative strength and places the burden of avoiding a stalemate on the

Figure 2.1: Game trees in *Transparency* and *Non-Transparency*



Note: The figure presents the game trees employed in the *Transparency* treatment (when disregarding the dotted-line information set of P) and in the *Non-Transparency* treatment (when including the dotted-line information set of P). As usual, payoffs are denoted at the end of the tree. The number at the top denotes the payoff of the proposer, while the number at the bottom is the payoff of the responder.

responder. In a sense investigating strategic ignorance on the proposer's side in a situation where she has already relative strength is conservative because one could assume that she would see a stronger need to improve her strength by remaining ignorant even more in a position of less relative strength.

Figure 1 presents the game tree with the exact payoff details. The acceptance of option A pays more to the proposer than the acceptance of option B, regardless of the state of nature. Whether, from the responder's perspective, the acceptance of option A is more preferable than the acceptance of option B depends on the actual state of nature. In state s_n , the responder's payoff from option B is higher than that from option A. The opposite is the case in state s_a . Thus, in state s_a , the payoffs of the two players are aligned, i.e., option A is both players' preferred option, while in state s_n , they are not aligned.

For comparability, we essentially use the same payoff parameters as Dana et al. (2007) in their dictator game.⁴ When the payoffs are aligned

⁴Note that the parameters used in Dana et al. (2007) put pressure on the proposer to inform herself because ignorantly choosing the selfish option A prevents implementing a (more) payoff-equalizing and efficient outcome in the case of state s_n . In contrast to Dana et al. (2007), we are not interested in a dictator game but in an ultimatum bargaining setting where the proposer relies on the acceptance decision of the responder. This puts additional pressure on the proposer to inform herself. To balance this pressure, we provide the proposer with a potential excuse to remain ignorant by slightly modifying the payoff structure used in their game. We reduce the option B payoff of the proposer by one unit. In Dana et al. (2007), in state s_n , both players earn 5 under option B, while in our setting, the proposer

(our state s_a), option A is more efficient in terms of maximizing total surplus, while option B is more efficient when the payoffs are not aligned (our state s_n). As in Dana et al. (2007), ex-ante, i.e., before the actual state is known, option A maximizes total expected payoffs.

Our five treatments build on this baseline game. In all of our treatments, the responder knows the actual state of nature when he decides whether to accept or reject the proposed offer. Treatments differ with respect to what the proposer knows or can learn about the actual state of nature before making the offer. We also vary the responder’s knowledge concerning what the proposer knew when making the offer. In the following, we introduce the details of our five treatments.

2.4.2 Treatments

In our first treatment, *Transparency*, we employ the game depicted in Figure 1 when disregarding the information set of the proposer indicated by the dotted line. The proposer knows the actual state of nature when she makes the offer. Our second treatment, *Non-Transparency*, employs the game that is described by the game tree in Figure 1 when including the dotted-line information set. The proposer is unaware of the true state of nature when making her offer. Regardless of the state of nature, however, it is beneficial for the proposer to offer option A , assuming that the responder accepts the offer. In the third treatment, *Choice*, we endogenize transparency, i.e., the proposer can choose between a transparent situation and a non-transparent one. The proposer can decide to inform herself about the actual state of nature or remain ignorant (both at no cost) before she decides on the offer. The responder accepts or rejects the offer after he learns whether the proposer informed herself about the actual state of nature, i.e., the responder is aware of whether an option A offer in state s_n was made knowingly or in the dark (the game trees of the games employed in this and the other two treatments can be found

earns one unit less than the responder. This allows the responder to more easily accept option A (which an ignorant proposer will offer) in state s_n because he might acknowledge that the proposer cannot be expected to offer an option where she earns less than the responder. Moreover, the sacrifice incurred by the proposer by proposing option B instead of option A is higher under our parametrization than in Dana et al.’s (2007) game.

in the appendix). Building on the game used in treatment *Choice*, we relax the assumption that the responder is informed of what the proposer knows when making the offer. In the fourth treatment, *Choice Uncertain Information Acquisition* (henceforth *Choice-UI*), it is uncertain whether the proposer will be successful in her attempt to inform herself about the true state of nature. If she chooses to remain ignorant, she remains ignorant for certain. If she chooses to inform herself, information acquisition is not certain, but there is a 50% chance that it will be successful. Otherwise, she remains ignorant. The responder is informed of whether the proposer knew the actual state of nature when she made the offer. If the proposer knew the state, it is clear to the responder that the proposer chose to inform herself and that her attempt was successful. If the responder learns that the proposer was not informed, it might be due to two different reasons. First, the proposer did not want to know. Or, second, she actually wanted to know but was not successful in informing herself. Thus, in the latter case, the responder cannot be sure about the actual intentions of the proposer, i.e., whether she tried to inform herself about the state. In our fifth treatment, *Choice Hidden* (henceforth *Choice-H*), the responder is kept uninformed about whether the proposer informed herself about the actual state of nature. This treatment enables us to disentangle two motives for remaining ignorant. One motive concerns self-image, i.e., wanting to be ignorant oneself (e.g., Bénabou and Tirole, 2006). The other motive is to signal to the responder that one is ignorant. In the treatments *Choice* and *Choice-UI*, these two motives cannot be separated.

In all treatments, the proposer knows what the responder will or will not learn about her chosen actions from the beginning.

2.4.3 Procedures

The experimental sessions took place at the Cologne Laboratory for Economic Research from August 2010 to April 2011. Subjects were recruited through the online recruiting system ORSEE programmed by Greiner (2003). We had a total of 592 participants (289 female) who were randomly drawn from a pool of over 3,000 students. Each session involved 16 to 32 participants who were not allowed to participate in more than one session. Approximately half of the participants were economics or

business administration majors; the other half were enrolled in different fields such as law and sciences. On average, participants were in their fourth year of study. We conducted two sessions of the *Transparency* ($n=64$), the *Non-Transparency* ($n=64$) and the *Choice-H* ($n=64$) treatments. As there are more potential variations of plays in the other two treatments, we ran six sessions of the *Choice* ($n=190$) and eight sessions of the *Choice-UI* ($n=210$) treatments.

At the beginning of a session, participants were randomly allocated to cubicles. After they had taken their seats, written instructions were distributed. Within a session, all subjects received the same instructions. The instructions for the different treatments were identical, with the exception of well-defined passages that described the treatment variations (see appendix).⁵ Each participant learned that he would play a simple one-shot game by interacting with one other person in the room who would be randomly and anonymously matched to him. We decided to employ a one-shot setting to focus on behavior that is not shaped by any type of endogenous social norm building, which could occur in a repeated setting. Additionally, by focusing on a one-shot setting, we attempted to avoid strategic considerations, e.g., to punish intentionally ignorant proposers in the early rounds by rejecting their offers to convince them not to remain ignorant in subsequent rounds. Before a session began, subjects had to complete a pen and paper quiz to confirm that they had understood the game (see appendix for the quiz).

The experiment was computerized using zTree software of Fischbacher (2007). Participants were informed of their randomly assigned roles on the first screen. We neutrally labeled a proposer as “participant X” and a responder as “participant Y”. The two states of nature were denoted “Case 1” and “Case 2”. In the treatments *Transparency* and *Non-Transparency*, a proposer simply had to choose between option *A* or option *B*. In addition, in the other three treatments, a proposer had to decide whether to inform herself about the true state of nature. As mentioned above, a responder always learned about the true state of nature before deciding whether to accept or reject the offer. To collect additional data on the responders’ behavior, we solicited their decisions using a re-

⁵The original instructions are written in German. The instructions provided in the appendix are translated from German into English.

duced version of the strategy method (Selten, 1967)⁶, i.e., the responder was not asked for a complete strategy but had to decide whether – given the actual state of nature – he would be willing to accept or reject option *A* and option *B*.

To gain deeper insights into subjects’ preferences and motivations, we asked them about their beliefs regarding the behavior of the other player and their decisions in hypothetical situations after the subjects had completed their decisions (and before they learned those of the other player). The hypothetical questions concerned a situation from a different treatment or a different state of nature. In the *Transparency* and *Non-Transparency* treatments, we asked the proposer and the responder to imagine that they had to decide while playing the same role but in the other treatment. In the *Choice* treatments, we asked subjects to imagine that the opposite decision had been made regarding the revelation of the true state of nature. For example, we asked a proposer who actually decided to remain ignorant which option, *A* or *B*, she would have offered, had she informed herself about the actual state. The responses to the hypothetical questions were not incentivized.

At the end of each session, the subjects were informed of the decisions of the other player they were matched with and their payoffs. They were then asked to complete a questionnaire on the motivations for their decisions. Finally, the participants privately received their payoffs from the game in addition to an individual participation fee of €2.50. On average, participants earned €7.04 (including the participation fee), and the sessions lasted for approximately 35 minutes.

2.5 Results

2.5.1 Results *Transparency* and *Non-Transparency*

In the first step, we will discuss the results from the treatments *Transparency* and *Non-Transparency*. Then, we will present our findings from the three *Choice* treatments in which subjects can deliberately remain

⁶There is no clear evidence on whether employing the strategy method leads to differences in behavior. On this “hot-or-cold” debate, see for example Brandts and Charness (2011). They find some differences in behavior in games where the second mover has a punishment or rejection choice.

Table 2.1: Results from *Transparency* and *Non-Transparency*

	Option A		Option B	
	State s_n	State s_a	State s_n	State s_a
<i>Transparency</i> ($n=64$)				
Proposals	9/16 (56%)	16/16 (100%)	7/16 (44%)	0/16 (0%)
Acceptance Rates	11/16 (69%)	16/16 (100%)	16/16 (100%)	10/16 (63%)
<i>Non-Transparency</i> ($n=64$)				
Proposals	30/32 (94%)		2/32 (6%)	
Acceptance Rates	15/16 (94%)	16/16 (100%)	16/16 (100%)	14/16 (88%)

Note: In the treatment *Non-Transparency*, the 32 proposers are unaware of the actual state.

ignorant.

Table 2.1 summarizes the behavior in the *Transparency* and *Non-Transparency* treatments. In *Transparency*, 9 of 16 proposers (56%) offered option A in state s_n .

In state s_a , all 16 proposers offered option A. In response to these proposals, 11 of 16 responders (69%) accepted option A in state s_n . As expected, the option B offers were always accepted (16/16, 100%) in state s_n . In state s_a , option A offers were always accepted (16/16, 100%), and option B offers were accepted by 10 of 16 proposers (63%).⁷

In *Non-Transparency*, all but two of the 32 proposers offered option A. A total of 15 of 16 responders (94%) accepted option A offers in state s_n , and 14 of 16 (88%) agreed to option B in state s_a . The option B offers in state s_n were always accepted. The same is true for option A offers in state s_a . In *Non-Transparency*, significantly more proposers offered option A in state s_n than proposers in *Transparency*.⁸

As hypothesized, responders reacted differently to option A offers in state s_n in the two treatments. In state s_n , responders accepted option A offers more frequently in *Non-Transparency* (94%) than in *Transparency* (69%). This difference is modestly significant.⁹

In the treatments *Transparency* and *Non-Transparency*, whether the proposer was informed about the actual state was exogenously determined, i.e., intentions about remaining ignorant or not did not play a

⁷Note that none of the proposers actually offered option B in state s_a , but because we employed a reduced version of the strategy method, responders were asked to provide their responses to both possible offers.

⁸With $p=0.019$ (Fisher-test, one-sided). This finding is also supported by OLS and probit regressions (see Table A2.2 in the appendix).

⁹With $p=0.086$ (Fischer-test, one-sided).

role. We now turn to the treatments where ignorance was endogenous, i.e., could be chosen by the proposer.

2.5.2 Results *Choice* Treatments

A non-trivial number of proposers decided to remain ignorant in the treatments *Choice* and *Choice-UI*. In both treatments, 24% of the proposers chose not to inform themselves of the actual state of nature. In *Choice-UI*, 53% of the proposers who attempted to inform themselves of the actual state were successful in acquiring this information, while the others remained ignorant. In *Choice-H*, 88% of the proposers informed themselves of the state.

Observation “Information Choices of Proposers”: In *Choice* and in *Choice-UI*, a considerable number of proposers chose to remain ignorant. In *Choice-H*, few proposers did not inform themselves of the state of nature.

Figure 2.2 displays the percentage of proposers who offered option *A* in state s_n . The behavior in state s_a - not shown in Figure 2.2 - was very similar to the behavior in *Transparency* and *Non-Transparency*: proposers nearly always offered option *A*, and responders nearly always accepted these offers (see the summary Table A2.1 in the appendix).


In accordance with our hypotheses, in each of the treatments, *Choice*, *Choice-UI* and *Choice-H*, very few proposers (in *Choice*, 3 of 35 proposers, in *Choice-UI*, 3 of 20 proposers and in *Choice-H*, 3 of 14 proposers) who (successfully) informed themselves of the actual state to be s_n , offered option *A*. The vast majority of proposers who informed themselves offered option *B* in this state. Proposers who remained ignorant almost always proposed option *A*.

10

Observation “Proposers’ Offers”: In all three *Choice* treatments, ignorant proposers proposed option *A* significantly more frequently than proposers who informed themselves of state s_n . A clear majority of proposers who learned that the state was s_n offered option *B*.

¹⁰In *Choice* with $p < 0.001$, in *Choice-UI* with $p < 0.001$ and in *Choice-H* with $p = 0.011$

Figure 2.2: Percentages of proposers offering option A in state s_n



Note: In *Non-Transparency*, *Choice ignorant*, *Choice-UI ignorant* and *Choice-H*, ignorant proposers are unaware that they are actually in state s_n ; the light gray shaded bars indicate the settings in which proposers knew the state, while the dark gray shaded ones indicate the settings in which they did not. For reasons of completeness, we included the *Choice-H ignorant*, although only four observations exist.


Proposers in the *Choice* treatments who chose to inform themselves of the state of nature were in a similar situation as proposers in the *Transparency* treatment when they made their offers. We found that proposers in *Choice* and *Choice-UI* who informed themselves that the state was s_n offered option B significantly more frequently than proposers who happened to be in state s_n in *Transparency*.¹¹

Figure 2.3 depicts the responders' rates of accepting option A in state s_n . In state s_a - not shown in Figure 2.3 - proposers always offered option A , which was nearly always accepted by the responders. In *Choice*, responders knew whether they received an offer from a proposer who informed herself of the state (*Choice informed*) or a proposer who chose to remain ignorant (*Choice ignorant*). In state s_n , 14 of 35 responders (40%) accepted option A when it was offered by a proposer who informed

(all two-sided Fisher-tests). See also the OLS and probit regressions in the appendix. There, the p -values show that compared to *Transparency*, a significantly smaller proportion of proposers who informed themselves offered option A . Proposers who remained ignorant offered option A significantly more often.

¹¹With $p < 0.001$ comparing *Transparency* and *Choice*, and $p = 0.014$ comparing *Transparency* and *Choice-UI* (Fisher-tests, two-sided).

Figure 2.3: Percentages of responders accepting option A in state s_n



Note: The light gray shaded bars indicate the settings in which proposers knew the state, the dark gray shaded ones indicate the settings in which they did not (except for *Choice-H*, where most proposers informed themselves of the state but responders were not informed of the proposers' decision of whether to inform themselves).

herself. In comparison, if option A was offered by ignorant proposers, these offers were accepted more frequently, i.e., by 7 of 12 responders (58%). Therefore, option A offers in state s_n from ignorant proposers seemed to be regarded as more acceptable than those from proposers who informed themselves of the state.

Responders in *Choice-UI* knew when a proposer successfully informed herself of the state of nature. In *Choice-UI*, 8 of 20 responders (40%) accepted option A when it was offered by proposers who successfully informed themselves that the state was s_n . If proposers in *Choice-UI* remained ignorant, responders did not know whether the proposers' ignorance was intentional or whether they unsuccessfully attempted to inform themselves of the state. We found that 21 of 31 responders (68%) accepted option A when it was offered by ignorant proposers. The difference in the acceptance rates for option A between proposers who informed themselves about the state and proposers who remained ignorant is significant.¹²

Thus, option A offers in state s_n offered by ignorant proposers were

¹²With $p=0.051$ (χ^2 -tests, two-sided).

accepted slightly more frequently in *Choice-UI* than in *Choice* (the difference is not significant). Option *A* offers in *Choice* and *Choice-UI* after deliberately informing oneself about the state s_n were disliked more by the responders than offering the option *A* in state s_n in *Transparency*.¹³

In *Choice-H*, the responder was not informed whether the proposer informed herself of the state of nature. In this treatment, the acceptance rate for option *A* in state s_n was between those of *Choice* informed and *Choice ignorant*, as 8 of 16 responders (50%) accepted this offer. The same is true when comparing the acceptance rates of *Choice-H* and *Choice-UI*.

Observation “Responders’ Acceptance Decisions”: *In Choice in state s_n , option A offers were accepted with a slightly higher frequency (although the difference is not significant) when they were offered by ignorant proposers than by proposers who informed themselves of the state. The corresponding comparison is significant in Choice-UI. In Choice-H, the acceptance rate of option A offers in s_n was between those of Choice informed and Choice ignorant.*

2.5.3 Proposer Payoffs

A crucial question is whether it pays for the proposer to be (strategically) ignorant. As we elicited responders’ acceptance behavior for both possible options in the actual experiment, the expected payoffs (henceforth *EP*) of different proposers’ strategies can be calculated.¹⁴ Table 2.2

¹³With $p=0.042$ (χ^2 -tests, two-sided, pooled data from *Choice* and *Choice-UI*). See also the OLS and probit regressions in Table A2.2 in the appendix. The p -values show that compared to *Transparency*, a significantly smaller share of responders accepted option *A* offers by proposers who informed themselves that the state was s_n .

¹⁴Under the matching strategy in our experiment, proposers who are kept ignorant in the *Non-Transparency* treatment earn significantly more than in the *Transparency* treatment, independent of the actual state ($p=0.016$, Mann-Whitney-U-test, henceforth: MWU, two-sided). In *Transparency*, proposers earn €5.19 on average compared to €5.69 in *Non-Transparency*. In the *Choice* treatment, there is no significant difference between the payoffs earned by informed and ignorant proposers ($p=0.158$, MWU, two-sided). However, with ambiguity about the proposer’s intentions in *Choice-UI*, ignorant proposers earn significantly more than proposers who successfully informed themselves that the state was s_n or s_a .

Table 2.2: Expected payoffs (EP) and observed total surplus

	EP (option A)		EP (option B)		Observed total surplus
	State s_n	State s_a	State s_n	State s_a	
<i>Transparency</i>	3.36	6	4	2.52	8.37
<i>Non-Transparency</i>	5.82		3.76		8.66
<i>Choice informed</i>	2.4	6	4	1.9	9.65
<i>Choice ignorant</i>	4.74		2.98		7.38
<i>Choice-UI informed</i>	2.4	5.7	4	2.52	8.65
<i>Choice-UI ignorant</i>	5.04		3.26		8
<i>Choice-H informed</i>	3	6	4	3	9.53
<i>Choice-H ignorant</i>	4.5		3.5		9

Notes: Expected payoffs (EP) for the proposers and observed total surplus for proposers and responders are displayed in Euros. The expected payoffs for the proposers are based on the observed responders' acceptance rates.

summarizes these EP for all possible proposer actions. Thus, in treatment *Transparency* under state s_n , it is beneficial for the proposer – in expected payoff terms – to offer option B , and in the case of state s_a , proposing option A is beneficial. In treatment *Non-Transparency*, where the proposer does not know the actual state, the expected payoff from offering option A is higher than that of offering option B . The best strategy for a self-regarding proposer who informs herself of the state in the three *Choice* treatments is to propose option B in the case of state s_n and option A in the case of state s_a . For a proposer who remains ignorant, the EP are always higher when offering option A than when offering option B . The EP of a proposer who informs herself in the treatment *Choice* and *Choice-H* are higher compared to a proposer who remains ignorant.¹⁵

In treatment *Choice-UI*, however, the expected payoff of a proposer who remains ignorant is slightly higher than that of a proposer who informs herself of the state. In addition, Table 2.2 also provides the actually observed total surplus for each treatment and information acquisition decision. In the treatments *Transparency* and *Non-Transparency*, the total surplus is nearly identical. In the three *Choice* treatments, however, the

($p=0.045$, MWU, two-sided). Remaining strategically ignorant in this treatment pays off, as proposers who inform themselves of the state earn €4.83 on average and ignorant proposers on average make €5.20.

¹⁵For example, in treatment *Choice*, the ex-ante expected payoff of a proposer who informs herself and optimally reacts to the observed responders' strategies would be €5 (in state s_n , she would choose option B , and in state s_a , she would choose option A), and the payoff would be €4.74 for a proposer who remains ignorant.

total surplus is always higher in situations where the proposers decided to inform themselves of the state than in situations where the proposer remained ignorant. Thus, more efficient outcomes are achieved when proposers inform themselves of the state.

2.5.4 Hypothetical Decisions and Beliefs in *Choice*

We asked ignorant proposers in *Choice* what they would hypothetically have done had they informed themselves of the state of nature. A total of 12 of the 23 ignorant proposers (52%) would have offered option *A* in state s_n . However, only 3 of the 35 proposers (9%) who actually informed themselves that the state was s_n offered option *A*. A potential explanation might be that proposers who informed themselves of the state of nature were more fairness-oriented than proposers who chose ignorance. It might be that the more fair-minded participants sorted into informing themselves (for a theoretical analysis related to this observation, see for example Grossman and van der Weele, 2013).

We also asked proposers who informed themselves of the state of nature whether they believed that option *A* in state s_n would have been accepted if offered by a proposer who informed herself. A total of 14 of the 71 proposers (20%) believed that such offers would have been accepted. The same proposers were also asked about their beliefs concerning the acceptance of an option *A* offer in state s_n if a proposer had chosen ignorance. Here, 26 of the 71 proposers (37%) who informed themselves believed that these offers would have been accepted. Posing the same questions to proposers who did not inform themselves of the state indicates that 10 of 23 (43%) believed that an option *A* offer in state s_n would have been accepted if the proposer had informed herself of the state of nature. Without informing themselves, 19 out of 23 (82%) ignorant proposers believed that such an offer would have been accepted.¹⁶ Thus, the proposers' beliefs seemed to reflect their different perceptions of the effectiveness of ignorance as a potential strategic advantage.¹⁷

¹⁶The differences in the beliefs of ignorant proposers and proposers who informed themselves concerning the effect of ignorance is significant ($p=0.027$, MWU).

¹⁷Asking participants to briefly explain the motivations behind their decisions generated interesting insights. A proposer who decided to inform herself of the state, for example, wrote (translated from German): "I informed myself of the state because I wanted to offer option *B* in the case of state s_n . If I really wanted to offer option *A* in this state, I would not have informed myself of the state to positively influence

Examining the responders' hypothetical decisions in the *Choice* treatment in greater detail sheds some light on the role of the proposers' intentions. Only 14 of the 35 responders (40%) actually accepted option *A* offers in state s_n offered by a proposer who informed herself. However, 21 of these 35 responders (60%) would have hypothetically accepted such offers if the proposers had remained ignorant.¹⁸ Moreover, 7 of 12 responders (58%) accepted option *A* offers in state s_n when they were offered by ignorant proposers. Only 4 of these 12 responders (33%), however, would have accepted such offers if they had been made by a proposer who informed herself of the state and found the state to be s_n .¹⁹

2.6 Concluding Remarks

The results from the *Transparency* and *Non-Transparency* treatments provide support for Schelling's conjecture: ignorance can indeed be an advantage for proposers. Nearly all uninformed proposers in the *Non-Transparency* treatment obtain their maximum payoff, i.e., €6. Responders seem to acknowledge that their opponents do not know the responders' payoff structure. The high rate of acceptance may be due to the responders' tendency to attribute the negative consequences of option *A* offers in state s_n to the randomly occurring state of nature. Thus, a selfish intention on the part of the proposer might be diluted.

In the three *Choice* treatments, there are at least two possible reasons

the responder to accept option *A*." Another proposer, who did not inform herself of the state of nature, commented: "I did not inform myself of the state to have an excuse for offering option *A* in state s_n . In my opinion, the responder then does not think that I am intentionally mean." A responder who received an offer from a proposer who learned the state wrote: "I accepted option *B* in state s_n , as it is the fair solution for both players. But I did not accept option *A* because I don't want to accept €6 for him and only €1 for me. I decided this way because I knew that the proposer informed herself of the state. In case the proposer would not have informed herself, I would have accepted all offers because then chance would have decided and the proposer would not have known what state actually occurred."

¹⁸Asking this question to all responders who received an offer from a proposer who informed herself regardless of whether they were in state s_n or state s_a showed that 38 of 71 responders (53%) would have accepted option *A* in state s_n from an ignorant proposer.

¹⁹Asking this question to all responders who received an offer from an ignorant proposer regardless of whether they were in state s_n or state s_a showed that only 8 of 23 responders (35%) would have accepted option *A* from a proposer who was aware that the state was s_n .

for a proposer to remain ignorant. The first is that a proposer wants to remain ignorant to herself, i.e., she does not want to know what a given offer means for a responder. A proposer who wants to offer option A may feel more morally consistent when ignorant compared to knowing the state, as the random draw could be blamed for an outcome that is potentially unfavorable for the responder (see Dana et al., 2007). However, our observations from treatment *Choice-H* indicate that this explanation seems to have limited force: when responders are not informed of the proposer’s decision of whether to acquire information, very few proposers remain ignorant (to themselves). A second reason for remaining ignorant might be that a proposer strategically uses ignorance. Such a proposer may believe that ignorance increases the responder’s inclination to accept option A offers in state s_n . This second explanation is supported by the observed proposers’ beliefs and responses to our open questions: proposers who deliberately remained ignorant believed that option A in state s_n would be less frequently accepted if they had informed themselves of the state.

There are, however, at least two reasons for the proposers to inform themselves of the state of nature. The first is that a proposer might propose option A in state s_n when she remains ignorant, which could lead to the responder’s rejection. Second, if a proposer has genuine pro-social preferences, she wants to acquire information about the true state of nature to offer option B in state s_n that gives the responder a higher payoff and is also more equal and efficient. This second explanation is supported by the finding that proposers who inform themselves of the state in the *Choice* treatments more frequently offer option B in state s_n than proposers in *Transparency*. Additionally, the hypothetical decisions also indicate that genuine pro-social proposers sort into informing themselves, whereas ignorant proposers attempt to strategically exploit the moral wiggle room.

At first sight, one might imagine that an option A offer in state s_n is evaluated similarly by a responder regardless of whether a proposer informed herself or deliberately remained ignorant. As it is costless to inform oneself of the state, it could, however, be argued that remaining intentionally ignorant is a more ruthless behavior on the part of the proposer. The acceptance rates show that responders tend to accept

option A offers in state s_n more frequently when they come from ignorant proposers than from proposers who informed themselves. This difference is even significant in *Choice-UI*, where responders could not be certain whether this ignorance was intended by the proposer. Therefore, Schelling’s conjecture (1960) that informational weakness can be a strength is supported (by the comparison of *Transparency* and *Non-Transparency*) but might also be extended: ignorance can even be used strategically if the opponent is aware of the ignorance but is uncertain whether the ignorance was intentional (shown in *Choice-UI*). Moreover, option A offers in state s_n from proposers who inform themselves of the state are more frequently rejected in each of the *Choice* treatments than in the *Transparency* treatment. Responders seem to perceive option A offers by proposers who deliberately informed themselves that the state was s_n (in the *Choice* treatments) as having worse intentions than proposers who offered option A and were immediately provided with the information that the state was s_n (in the *Transparency* treatment).

In light of our results, insufficient attention has been devoted to research on the strategic use of ignorance in bargaining. We are aware that our experiment uses a specific bargaining format and a specific payoff structure. Further research is needed to verify whether our findings extend to richer bargaining formats that go beyond take-it-or-leave-it bargaining, for example, sequential offer bargaining. In such a situation, the bargaining power is more equally divided between the two players, and therefore one could expect that a proposer has a greater need to strengthen her position, for example, by exploiting strategic ignorance. As we have a situation characterized by asymmetric information where the responder knows the state of nature but the proposer does not, sequential offer bargaining could involve signaling on the side of the responder, which makes the strategic interaction more complex. Bargaining settings are often of a repeated nature, and one might wonder whether parties adapt their behavior over time. Thus, it would also be interesting to study learning and endogenous social norm building in repeated settings. Because remaining ignorant might be perceived as rude behavior, it is an open question as to whether strategic ignorance will be used less frequently in a face-to-face environment.

2.7 Appendix

Figure A2.1: Game-tree of the *Choice* treatment

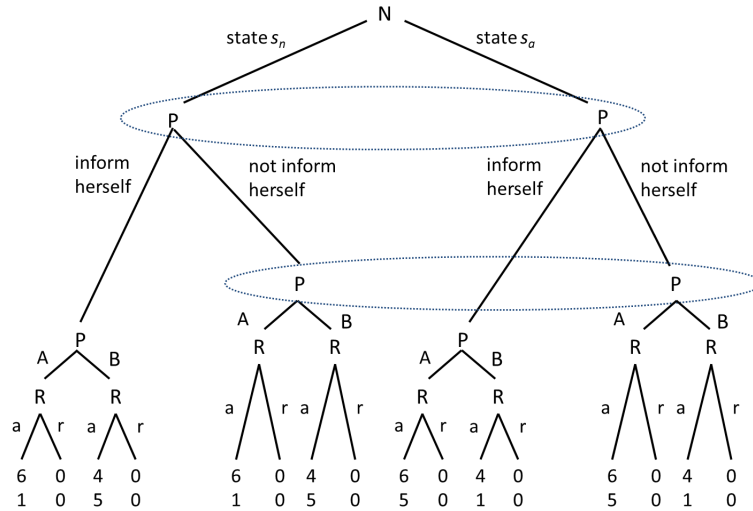


Figure A2.2: Game-tree of the *Choice Uncertain Information Acquisition (Choice-UI)* treatment

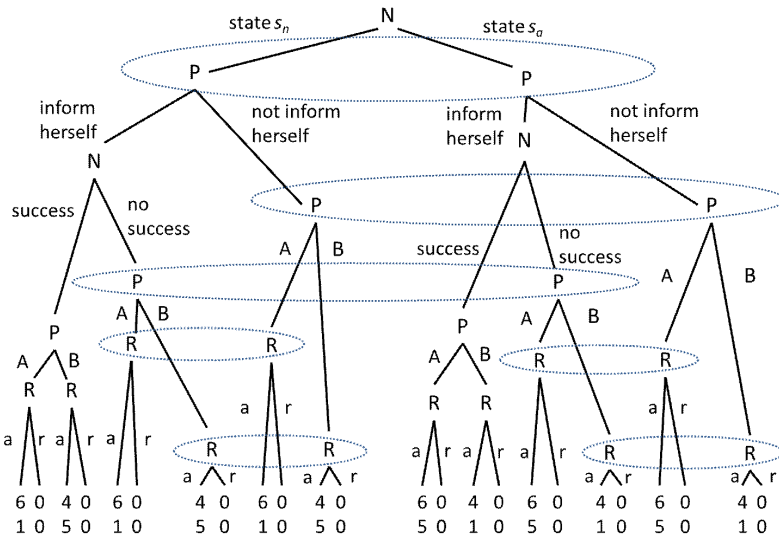


Figure A2.3: Game-tree of the *Choice Hidden* (*Choice-H*) treatment




Table A2.1: Proposers' offers and receivers' acceptance rates

Treatment	Proposals / Acceptance rates	Transparent / Informed				Non-Transparent / Not Informed			
		Option A		Option B		Option A		Option B	
		State s_n	State s_a	State s_n	State s_a	State s_n	State s_a	State s_n	State s_a
<i>Transparency</i>	Proposer	9/16 (56%)	16/16 (100%)	7/16 (44%)	0/16 (0%)	-	-	-	-
	Responder	11/16 (69%)	16/16 (100%)	16/16 (100%)	10/16 (63%)	-	-	-	-
<i>Non-Transparency</i>	Proposer	-	-	-	-	30/32 (94%)	2/32 (6%)	2/32 (6%)	-
	Responder	-	-	-	-	15/16 (94%)	16/16 (100%)	16/16 (100%)	14/16 (88%)
<i>Choice</i>	Proposer	3/35 (9%)	37/37 (100%)	32/35 (91%)	0/37 (0%)	23/23 (100%)	0/23 (0%)	0/23 (0%)	-
	Responder	14/35 (40%)	37/37 (100%)	35/35 (100%)	18/37 (49%)	7/12 (58%)	12/12 (100%)	11/11 (100%)	6/11 (55%)
<i>Choice-UI</i>	Proposer (intentionally ignorant)	3/20 (15%)	22/22 (100%)	17/20 (85%)	0/22 (0%)	24/25 (96%)	1/15 (4%)	1/15 (4%)	-
	Proposer (unintentionally ignorant)	8/20 (40%)	21/22 (95%)	20/20 (100%)	14/22 (63%)	35/38 (92%)	3/38 (8%)	3/38 (8%)	-
<i>Choice-H</i>	Proposer	3/14 (21%)	14/14 (100%)	11/14 (79%)	0/14 (0%)	4/4 (100%)	0/4 (0%)	0/4 (0%)	-
	Responder	-	-	-	-	8/16 (50%)	16/16 (100%)	16/16 (100%)	12/16 (75%)

Notes: Proposers in the treatment *Non-Transparency* and proposers who did not inform themselves about the state in the *Choice* treatments are not aware of the actual state. In the treatment *Choice-H* the responders do not know whether the proposers informed themselves about the state or not.

Table A2.2: OLS and probit regressions on decisions in state s_n

	(1) OLS	(2) Probit	(3) OLS	(4) Probit
Dependent variable:	<i>offer_A</i>	<i>offer_A</i>	<i>accept_A</i>	<i>accept_A</i>
<i>Non-Transparency</i>	0.375*** (0.118)	0.470*** (0.134)	0.249 (0.155)	0.320** (0.126)
<i>Choice_i</i>	-0.477*** (0.101)	-0.510*** (0.104)	-0.288** (0.129)	-0.301** (0.134)
<i>Choice_ni</i>	0.437*** (0.127)	dropped	-0.105 (0.169)	-0.123 (0.184)
<i>Choice-UI_i</i>	-0.413*** (0.112)	-0.401*** (0.114)	-0.289** (0.116)	-0.310** (0.123)
<i>Choice-UI_i_ns</i>	0.326*** (0.115)	0.389*** (0.148)		
<i>Choice-UI_ni</i>	0.361*** (0.124)	0.439*** (0.146)	0.157 (0.115)	0.164 (0.112)
<i>Choice-H_i</i>	-0.348*** (0.122)	-0.330** (0.134)		
<i>Choice-H_ni</i>	0.437* (0.250)	dropped		
<i>Choice-H</i>			-0.188 (0.155)	-0.206 (0.165)
Constant	0.563*** (0.0833)		0.688*** (0.101)	
Observations	146	132	146	146
<i>R</i> -squared	0.583		0.141	
pseudo <i>R</i> -squared		0.446		0.117

Note: Model (1) and (2) explain the effects of the different treatment situations on the proposer's decision to offer option A (*offer_A*). Model (3) and (4) explain the effect of the different treatment situations on the responder's decision to accept option A (*accept_A*). Reference category is the *Transparency* treatment. Independent variables are dummies for different treatment situations. *Choice_i* stands for the *Choice* treatment in which a proposer informed herself to be in state s_n and *Choice_ni* stands for a situation in which a proposer did not inform herself about the state. The same applies for the treatments *Choice-UI* and *Choice-H*. *Choice-UI_i_ns* stands for a situation, in which a proposer tried to inform herself about the state but was not successful, i.e., she remained uninformed about the state. *Choice-H* means a situation where a responder does not know if a proposer informed herself about the state or not. Standard errors are in parentheses. Levels of significance are indicated as follows: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Instructions of the Experiment (translated from German)

{[*Transparency, Non-Transparency, Choice, Choice-UI, Choice-H*]

Instructions

Welcome and thank for your participation in today's experiment. Please read these instructions carefully. If you have any questions, do not hesitate to raise your hand, we will help you personally.

In this experiment you can earn money. The amount of money you will earn depends on both, your own decisions and the decisions taken by another participant. At the end of the experiment you will receive your payoff in cash. Your payoff is composed out of the experiment's payoff and a show-up fee of € 2.50. You will receive the participation show-up fee independently from the payoffs you gain during the experiment. From now on, we kindly asked you stop all communication. Please make sure that your cell phone is switched off. A violation against these rules may lead to the exclusion from this and other experiments.

The Basic Decision Situation

In the decision situation there are two types of participants, participant X and participant Y, and two possible cases, case 1 and case 2. The payoff-tables for both participants are as follows:

Option A	X: 6 €	Y: 1€	Option A	X: 6 €	Y: 5€
Option B	X: 4 €	X: 5€	Option B	X: 4 €	X: 1€
Case 1			Case 2		

One of the two cases will randomly appear. Case 1 and case 2 occur with equal probability. Within each case, two Options exist, Option A and Option B. Participant X chooses one of the two Options. With each respective Option, different payoffs for participant X and participant Y are associated. Independently from the occurring case, Option A leads to a payoff of € 6 and Option B to € 4 for participant X. For participant Y different payoffs are associated with the different Options in the different cases. In case of case 1, Option A would lead to a payoff of € 1 and Option B to a payoff of € 5. In the event of case 2, Option A would lead to a payoff of € 5 and Option B to a payoff of € 1. Participant Y is able to accept or to reject the Option proposed by participant X. If participant Y accepts the Option chosen by participant X, both participants receive the respective payoffs. If participants Y rejects the Option chosen by participant X, both participants receive zero payoffs.

The Decision Situation in Detail

At the beginning of the experiment you will be informed via the computer screen which of the two possible roles – either participant X or participant Y – will be randomly assigned to you. At the same time and again randomly you will be assigned to another participant you will interact with. This assignment is completely anonymous. In each composed pair, one participant has the role of participant X and the other has the role of participant Y. The interaction within each pair only occurs through the computers. After you are informed about your role, you can continue by pressing an OK-Button.}

{[*Transparency*] Participant X takes the first decision. He is informed which of the two possible cases – case 1 or case 2 – has occurred. Thus, he sees one of the two following payoff-tables:}

{[*Choice, Choice-UI, Choice-H*] Participant X takes the first decisions. At the beginning he sees the following buttons. He chooses between the two buttons by pressing one of them:

Inform yourself about participant Y's payoff	Do not inform yourself about participant Y's payoff
---	--

Participant X has two possibilities: He can {[*Choice-UI: try to*]} inform herself about participant Y's possible payoffs or he cannot inform herself, respectively. If participant X decides for pressing the button “Inform yourself about participant Y's payoff” he sees {[*Choice-UI: with a probability of 50%*]} one of the two possible payoff-tables:

Option A	X: 6 €	Y: 1€	Option A	X: 6 €	Y: 5€
Option B	X: 4 €	X: 5€	Option B	X: 4 €	X: 1€
Case 1			Case 2		

Specifically, through pressing the button “Inform yourself about participant Y's payoff”, participant X is informed {[*Choice-UI: with a probability of 50%*]} which of the two possible cases – case 1 or case 2 – has occurred. {[*Choice-UI: With a probability of 50%*] participant X does not learn which case has occurred although he pressed the Button “Inform yourself about participant Y's payoff”. In this case the attempt to inform herself about participant Y's payoff was not successful and participant X only sees the following table:

Option A	X: 6 €	Y: ? €
Option B	X: 4 €	X: ? €
Case 1 or case 2		

Here participant X does not know which of the two possible cases has occurred. Instead of participant Y's possible payoffs only a “?” is visible.} {[*Transparency, Choice, Choice-UI, Choice-H*] Now, participant X chooses between Option A and Option B. Note that the payoffs for participant X associated with the two cases are identical with respect to the Options. This means participant X always receives €6 for Option A and €4 for Option B provided participant Y accepts the chosen Option. For participant Y the payoffs in both cases are differently.}

{[*Non-Transparency*] Participant X takes the first decision. He is not informed which of the two possible cases – case 1 or case 2 – have occurred. Thus, he sees the two following payoff-table:}

{[*Choice, Choice-UI, Choice-H*] If participant X decides for pressing the button “Do not inform yourself about participant Y's payoff” he sees the following payoff-table:

Option A	X: 6 €	Y: ? €
Option B	X: 4 €	X: ? €

Case 1 or case 2

Participant X then decides not to know the possible payoffs for participant Y.} {[*Non-Transparency, Choice, Choice-UI, Choice-H*] Specifically he is not informed if case 1 or case 2 occurred. Instead of participant Y's possible payoffs only a “?” is visible. But participant X knows that he receives a payoff of €6 for Option A and a payoff of €4 for Option B if participant Y accepts his decision. Participant X then chooses between Option A and Option B.}

{[*Transparency, Non-Transparency, Choice, Choice-UI, Choice-H*] Participant Y will be informed which of the cases – case 1 or case 2 – have occurred.} {[*Choice*] At the same time he is {[*Choice-H*] not} informed if participant X informed herself about participant Y's possible payoffs or if he covered it. Thus, he will be {[*Choice-H*] not} informed if participant X knew the possible payoffs for participant Y associated with the Option chosen. {[*Choice-UI*] If participant X was uninformed about participant Y's payoff this can be due to two reasons: either participant was not successful in informing herself about participant Y's possible payoffs or participant X decided not to inform herself about participant Y's possible payoffs. If participant X was uninformed participant Y will not be informed due to which reason participant X became uninformed.}

{ [*Transparency*, *Non-Transparency*, *Choice*, *Choice-UI*, *Choice-H*] Participant Y now decides for both possible Options of participant X – so for Option A and Option B – if he accepts or rejects participant X’s choice. Depending on the Option actually chosen by participant X, payoffs for both participants are determined. If participant Y accepts the Option chosen by participant X, both participants receive their respective payoffs. If participant Y rejects the Option chosen, both participants receive zero payoffs.}

{ } = Indicate the phrase that is exclusively employed in the respective treatment, [*Transparency*], [*Non-Transparency*], [*Choice*], [*Choice-UI*] or [*Choice-H*].

3 Lying and Team Incentives

3.1 Introduction

Deception and lying is common in all kinds of social interaction (Bok, 1999; Feldman, 2009) and recently, this topic also raised considerable interest in the experimental economics literature (Croson, 2005). In a simple and innovative die-rolling game Fischbacher and Föllmi-Heusi (2013), henceforth F&H, for example, find that people systematically over-report the true value of a private die-roll when the reported number determines their individual pay. Interestingly, people seem not to exaggerate their claims to the full extent what the authors call ‘incomplete lying’. This result is in line with the idea of ‘self-concept maintenance’ investigated by Mazar et al. (2008). They argue that lying to a small extent does not necessarily require changing one’s self-image as an honest person. Gneezy (2005) investigates the role of consequences on the inclination of lying and finds that people deceive more often the higher the own profit from lying and the lower the loss for the deceived person (for an extension see Sutter, 2009). Schweitzer and Hsee (2002) point out that people tend to justify lying more easily when other people benefit from dishonest behavior. Following this argument, Wiltermuth (2011) shows that people are more likely to cheat when the benefits of doing so are split with another person even if this other person is totally unknown to the cheater. In a similar vein, Mazar and Aggarwal (2011) demonstrate that individuals who are primed in a collectivist mindset more frequently behave unethical, i.e., offering bribes, as they feel less responsible for their own actions. Thus, deceit also seems to be psychologically easier because of diffusion and displacement of responsibility, i.e., the likelihood to deceive increases when individual causal agency for unethical behavior is obscured (Bandura, 1991; Bandura et al., 1996).

In this paper we take an organizational perspective and look at the in-

fluence of different compensation schemes on deceptive behavior and lying. Such a perspective is also taken by Cadsby et al. (2010) who employ an anagram task and experimentally analyze the differences in cheating under piece-rate, target-based and tournament incentive schemes. They find that lying in terms of over-claimed words is most pronounced under the target-based system as targets seem to encourage people to lie particularly if one is close to the target. By following up on their work, in the current study we compare lying behavior under the two probably most commonly observed incentive schemes in organizations, i.e., individual piece-rate compensation and team compensation schemes (Gibbons, 1998; Lazear and Gibbs, 2009). Ledford et al. (1995), for example, show that more than 70% of major US firms use some form of team-based rewards. Using data of a representative survey of German companies Berger et al. (2010) find that performance-related pay such as piece-rate compensation and team-based variable compensation are widespread. This suggests that the investigation of potential deceptive behavior under these two compensation schemes is important.

We use a variant of the die-rolling game of F&H, which in their version resembles an individual piece-rate compensation scheme. We are able to confirm F&H's findings, i.e., people systematically lie but quite often they are reluctant to do so to the full extent. Our team compensation scheme is modeled as a revenue sharing mechanism (for an earlier experimental study on revenue sharing see Nalbantian and Schotter, 1997). The production outputs of two agents are pooled and each agent receives one half of a compensation unit for each unit of the joint production output. Comparing the marginal incentives to lie under the two schemes reveals that under the team compensation scheme the marginal gain from lying, i.e., the return from exaggerating the own production output by one unit, is only half of the marginal gain from lying under the individual piece-rate scheme. Assuming increasing marginal costs of lying this could lead one to assume that lying should be more pronounced under the individual piece-rate scheme than under the team incentive scheme. On the other hand, in the team incentive scheme lying is not exclusively beneficial for oneself - as it is the case under the individual compensation scheme - but it also benefits the other agent in the team. Thus, an agent under a team incentive scheme might be more able to justify such a *white*

lie to herself compared to a lie under the individual scheme - after all she is doing something ‘good’ for the other team member (see Wiltermuth, 2011). Moreover, lying under the team incentive scheme might be psychologically easier because this compensation scheme promotes *diffusion of one’s own responsibility*, i.e., tracing lies back to individual team members is more difficult under team compensation compared to individual compensation.

Extending upon our two main experimental treatments we additionally run a control treatment in order to explore potential motives of lying behavior under the two compensation schemes. Disentangling different motives behind deceptive actions would shed some light on the effectiveness of intra-organizational arrangements to reduce lying behavior when distinct compensation incentives are present.

The discussed lines of reasoning point into different directions whether lying is more severe under an individual than a team incentive scheme. Our experimental results suggest that lying is in fact more pronounced under the team incentives than under the individual piece-rate scheme. We also find *diffusion of responsibility* to be a stronger driving force for lying in teams than the *white lie* justification.

3.2 Experiment

Our experiment employed a simple one shot decision task closely resembling the baseline treatment of F&H. We ran two waves of experimental sessions. The first wave included our two main experimental treatments. Within each session we randomly assigned subjects either to the piece-rate or the team based compensation scheme. The second wave consisted of a control treatment designed to investigate potential motives of lying behavior under each of the above compensation schemes. Due to the short nature of the task we followed the procedure of F&H in conducting the experiment after several different other experiments. Experimental sessions were run in the laboratories of Bonn University and the University of Cologne between August 2010 and August 2012 and involved 554 subjects (with a mean age of 24.55 and 51% being female).

At the end of each experiment we asked subjects to fill in a questionnaire for a statistical survey for which they would be rewarded indepen-

dently of the preceding experiment. The questionnaire contained questions about gender, age, and personality measured by a 10-item Big-Five inventory covering the traits Openness, Conscientiousness, Extraversion, Agreeableness and Neuroticism (Rammstedt and John, 2007).¹ A separate instruction sheet explained that their pay would be based on ‘points’ p_i that were randomly determined, i.e., by rolling a standard 6-sided die. By introducing points that were interpreted as ‘random production output’ we slightly adapted the setting of F&H. The reason was that we were particularly interested in investigating lying under different compensation schemes.

In our two main experimental treatments it was explained to the subjects that the points p_i of subject i resulted from the number d_i shown on the die, i.e., $p_i = d_i$, if $d_i \in \{1, \dots, 5\}$. If a 6 was diced ($d_i = 6$), no points were earned ($p_i = 0$). Subjects were randomly assigned to the two treatments that differed in the way points, i.e., random production outputs, were translated into payoffs π_i . The first main experimental treatment *Individual* closely resembled the baseline treatment in F&H, i.e., subjects were instructed that the payoff of agent i would be defined as $\pi_i = p_i$. In the second main experimental treatment *Team* a subject i was randomly and anonymously matched with a different subject j to form a team. Subjects were told that team-member i ’s individual payoff would be defined according to the following sharing rule: $\pi_i = 1/2 \cdot (p_i + p_j)$. Subjects also learned that team-member j ’s payoff would be exactly the same, i.e., $\pi_j = \pi_i$.

In our control treatment, *Team-Mixed*, we varied the subjects’ influence on the team members’ payoff. In this treatment two types existed, a player i with individual compensation and a player j with team compensation. Subject i was randomly and anonymously matched with a different subject j to form a team. Subject i was informed that her own payoff was determined according to the following rule: $\pi_i = 1/2 \cdot (p_i + p_c)$ with p_c representing the output determined by a random die roll executed by the computer. It is important to note that the setting in *Team* differed from the individual compensation setting in *Team-Mixed*. In the

¹We decided to include a rather abstract but in HR-departments frequently used personality construct measure, which was apparently unrelated to the preceding die rolling task. In order to not induce an experimenter demand effect we refrained from explicitly asking about underlying motives regarding subjects’ behavior in the die rolling task.

latter subject i 's payoff did not depend on subject j 's reported outcome. The payoff of player j with team compensation in *Team-Mixed* was determined according to the following rule: $\pi_j = 1/2 \cdot (p_i + p_j)$. This payoff rule is similar to the one applied in our treatment *Team*, however, subject j in *Team-Mixed* did not produce a positive externality for subject i . Hence, the payoff of player i with individual compensation was determined by i 's own die roll and a random draw. The payoff of player j with team compensation was determined by j 's own die roll and the roll of the player with individual compensation. Both team members were also informed about the compensation rule of the respective other team member. We designed the *Team-Mixed* treatment such that the marginal gains from lying for both player types were the same as the marginal gains from lying in the *Team* treatment.

Subjects were asked to privately role the die in their cubicles and to jot down the appearing number on a sheet of paper, which was handed out to the subjects and collected afterwards.² This procedure ensured that the experimenter was not able to observe the truly diced numbers whatsoever and this was known to the subjects. Hence subjects could easily lie about their rolled number and consequently could secretly manipulate their payoffs that were solely dependent on their reported production output. After subjects had reported their production outputs they filled in the personality questionnaire which was administered after the die roll. At the end of the sessions, each subject k was privately paid π_k , which was calculated based on the reported production outputs and the respective payoff rule. Each payoff unit was worth €1. It should be emphasized that we designed the incentive schemes such that they were comparable with respect to two important characteristics. First, the expected payoff of a subject was €2.50 under all experimental treatments if one assumes that all subjects honestly reported their true numbers. Secondly, in our two main treatments *Individual* and *Team*, if all subjects behaved selfishly

²To be more precise we asked them to jot down the very first number that appeared on the die. In fact we followed the procedure of F&H and explicitly allowed the subjects to roll the die several times. As F&H we argued that by doing so subjects could assure themselves of the die being fair.

and maximally exaggerated their numbers they earned €5.

3.3 Hypotheses

We are primarily interested in whether agents are inclined to lie more under the individual piece rate or the team incentive scheme. As mentioned earlier two competing hypothesis are at hand. The first one relates to the fact that the marginal gain from lying is higher under the individual piece-rate scheme than under the team incentives. If one assumes increasing costs of lying one can derive our first hypothesis.

***Hypothesis 1:** More lying is observed under the individual piece-rate incentive scheme than under the team incentive scheme.*

On the other hand, under the team incentive scheme a liar has the excuse that lying comes with doing something good for the other member of the team, i.e., a lie is partly a *white lie*. Such an excuse is not available under the individual piece-rate incentive scheme. Furthermore, under the team incentive scheme subjects' individual deceptive acts are not unambiguously attributable to them individually. Due to this possible diffusion of responsibility of a seemingly overreported payoff subjects might be more inclined to overreport. These considerations result in our second hypothesis.

***Hypothesis 2:** More lying is observed under the team incentive scheme than under the individual piece-rate incentive scheme.*


To see which of the two hypotheses can be supported we have to turn to our data.

3.4 Results

Our main results are summarized in Figure 3.1 and in Table 3.1. In the figure we see the distribution of reported production outputs in our treatments. Rolling fair dices should generate a uniform distribution of production outputs (dashed line), i.e., every possible production output should come up with the same probability of $\frac{1}{6}$, and an average actual production output of 2.5.

The left side of Figure 3.1 shows the results of our treatment *Individual*. For comparative reasons the white bars indicate the results of the

Figure 3.1: Relative frequency of reported production outputs in the three different treatments



Note: The dashed line represents the expected true value of 0.16.

baseline treatment of F&H. Visual inspection already reveals that the results of our treatment *Individual* are very similar to their results. Indeed, a Mann-Whitney-U-test (MWU-test) comparing the distribution of reported production outputs between both treatments reveals that there is no statistically significant difference ($p=0.224$, two-sided). A comparison of the treatment *Individual* with the treatment *Team* - shown in the second panel of Figure 3.1 - provides insights regarding our research question.

Observation 1: *Reported production outputs are significantly higher in treatment Team than in treatment Individual.*

While subjects report on average 3.31 in treatment *Individual*, they report on average 3.86 in the treatment *Team*. Subjects report significantly higher production outputs in the treatment *Team* than in treatment *Individual* (MWU-test: $p=0.003$, two-sided).³ The observation that subjects

³F&H report an average of 3.52 in their baseline treatment. Also the distributions of numbers in F&H's baseline treatment and in *Team* are significantly different (MWU-test: $p=0.022$, two-sided).

Table 3.1: Descriptive statistics

Treatment	n	Type	\bar{p}_i	Reported production output p_i (rel. frequencies)					
				0	1	2	3	4	5
<i>Team</i>	132		3.86	.04 ⁻⁻⁻⁻	.05 ⁻⁻⁻⁻	.10 ⁻⁻	.08 ⁻⁻⁻⁻	.27 ⁺⁺⁺	.45 ⁺⁺⁺
<i>Individual</i>	156		3.31 ^{***}	.08 ⁻⁻⁻⁻	.12 ⁻	.08 ⁻⁻⁻⁻	.13	.26 ⁺⁺⁺	.31 ⁺⁺⁺
<i>Team-Mixed</i>	130	(I)	3.36 ^{***}	.10 ⁻⁻	.05 ⁻⁻⁻⁻	.10 ⁻⁻	.18	.27 ⁺⁺⁺	.30 ⁺⁺⁺
	136	(T)	3.63	.04 ⁻⁻⁻⁻	.10 ⁻⁻	.11 ⁻⁻	.12 ⁻⁻	.24 ⁺⁺	.40 ⁺⁺⁺
F&H	389		3.52 ^{**}	.06 ⁻⁻⁻⁻	.07 ⁻⁻⁻⁻	.12 ⁻⁻⁻⁻	.13 ⁻⁻	.27 ⁺⁺⁺	.35 ⁺⁺⁺

Notes: \bar{p}_i is the average reported production output. (I) represents the player with individual compensation in *Team-Mixed*, (T) stands for the player with team compensation in *Team-Mixed*. Stars show the significance of a two-sided Mann-Whitney-U test (*=10 %-level, **=5 %-level, ***=1 %-level) comparing the distribution of reported production outputs with the treatment *Team* with the hypothesis that reported production outputs are equally distributed. Plus and minus signs display the significance of a one-sided binomial test indicating that the observed relative frequency is smaller (larger) than $\frac{1}{6}$ (-(+)=10 %-level, --(++)=5 %-level, ---(+++)=1 %-level).

are more inclined to lie under the team incentive scheme is also supported, when comparing the frequencies of the maximal possible exaggeration, i.e., a reported production output of 5.

Observation 2: *A production output of 5 is reported more often in treatment Team than in Treatment Individual.*

This observation is backed by a χ^2 -test ($p=0.014$, two-sided). In line with the findings of F&H we also observe that lying is ‘incomplete’ in our two treatments.

Observation 3: *Incomplete lying is observed in both treatments, Team and Individual.*

Support for this observation is shown in Table 3.1. The results indicate that in both treatments the frequencies of reported production of 4 significantly exceed the frequency that one would expect from honest subjects. While in both treatments 0, 1 and 2 are reported significantly less often, only the frequency of the production output of 3 under individual incentives cannot be distinguished from the relative frequency of $\frac{1}{6}$.

Observation 4: *When investigating why lying is more pronounced in treatment Team than in treatment Individual the argument of “diffusion of responsibility” is likely to have a stronger force than the “white lie” justification.*

To generate deeper insights into the underlying motives for the more pronounced lying in *Team* than in *Individual* we run the control treatment *Team-Mixed*. One explanation for the increased lying in *Team* might be the justification to do something good for the other team member. To investigate this white lie justification we compare the players with team-compensation in *Team-Mixed* with the players in *Team*. Note that the reported production output of the player with team-compensation in *Team-Mixed* has no impact on the other team member. Otherwise the two incentive situations are the same.

The average reported production output of players with team compensation in the *Team-Mixed* is statistically indistinguishable from players in *Team* (3.63 vs. 3.86, MWU-test: $p=0.186$, two-sided). Since the output of the former compensation scheme has no effect on the earnings of the other team member one would expect that players in the treatment *Team* report higher outputs (compare our hypothesis 2). Thus, it is not very likely that the pronounced lying in the treatment *Team* in comparison to treatment *Individual* is mainly driven by the white lie justification.⁴

Another explanation for the differences in lying in *Individual* and *Team* could be diffusion of responsibility. To test this we compare behavior of players with individual compensation in *Team-Mixed* and behavior of players in *Team*. In the latter treatment the responsibility for possible deceptive actions might be obscured.

The players with individual compensation in *Team-Mixed* report on average a significantly smaller number than subjects in *Team* (3.36 vs. 3.86, MWU-test: $p=0.004$, two-sided). In both settings, reported out-

⁴Interestingly, our finding is different from the results reported by Wiltermuth (2011) who finds in a real-effort anagram task that people are more likely to cheat when another person also benefits from it. In line with our study, Danilov et al. (2013) investigate advice giving in an experiment with financial professionals. In general, they find no difference in deceptive self-serving advice giving when professionals operate under team versus individual incentives. Only when professionals in a team have established social ties they tend to convey more self-serving advice which harms clients.

puts have a positive effect on the other team members' earnings. The difference is that in *Team* it is less obvious who of the two team members actually lied. This is not the case for the player with individual compensation in *Team-Mixed*. Thus, diffusion of responsibility seems to encourage players in *Team* to lie more than players in *Individual*.⁵ In the following we will check the robustness of our main findings controlling for individual difference variables collected after the die roll task. Relating gender, age, and personality traits with (over)reported production outputs unveils some further interesting insights about potential determinants of lying behavior. In Table 3.2 we run a series of linear regression models to predict reported production outputs by stepwise including Treatments, Female, Age, and Big-Five personality factors as explanatory variables. In all three regression models treatment *Team* serves as a reference group. Models (1)-(3) show that our main finding on the influence of team incentives on lying behavior (Observation 1) is quite robust. Furthermore, these models underline that players under both individual incentive schemes report smaller numbers than players under the two team incentives schemes. Thus, diffusion of responsibility seems to encourage players in *Team* to lie more than players in *Individual*. Our finding is in line with Mazar and Aggarwal (2011) who have shown that in an interdependent setting, individuals feel less responsible for their own actions and this mediates the likelihood to transgress.

⁵The behavior of the player with individual compensation in *Team-Mixed* offers another valuable insight: One might argue that subjects have a minimum earnings goal that they want to reach by over-reporting production outputs and since the marginal gain from lying in *Team* is only half of that in *Individual*, cheating might be higher in *Team*. To test this claim we compare the reported production outputs of players with individual compensation from *Team-Mixed* (marginal gain equal to $\frac{1}{2}$) with those from *Individual* (marginal gain equal to 1). Note that in both treatments players cannot diffuse their responsibility, i.e., they do not benefit from another player's production output. We find that players with individual compensation from *Team-Mixed* (3.36) do not report significantly higher outputs than those from *Individual* (3.31) (MWU-test: $p=0.959$, two-sided). This is remarkable given the fact that players with individual compensation from *Team-Mixed* even have an excuse to over-report because another player would benefit from their lies. Thus, the motive of a minimum earnings goal seems not to play a decisive role.

Table 3.2: Explaining reported production output p_i

Dependent variable:	(1)	(2)	(3)
	Reported production output p_i		
<i>Individual</i>	-0.550*** (0.18)	-0.502*** (0.18)	-0.518*** (0.18)
<i>Team-Mixed (I)</i>	-0.502*** (0.19)	-0.351* (0.20)	-0.369* (0.20)
<i>Team-Mixed (T)</i>	-0.239 (0.18)	-0.165 (0.18)	-0.179 (0.18)
Female		-0.392*** (0.14)	-0.475*** (0.16)
Age		-0.0264** (0.01)	-0.0157 (0.01)
Openness			0.0616 (0.07)
Conscientiousness			-0.115 (0.09)
Extraversion			0.183** (0.07)
Neuroticism			0.177** (0.08)
Constant	3.864*** (0.12)	4.654*** (0.35)	3.646*** (0.59)
Observations	554	546	544
<i>R</i> -squared	0.020	0.039	0.059

Notes: OLS-regression coefficients (robust standard errors in parentheses), reference group: *Team* treatment; we did not include the personality factor Agreeableness in the regressions since its scale reliability was too low. The number of observation differs due to missing values. *Individual*, *Team-Mixed (I)* and *Team-Mixed (T)* represent dummy variables for the respective treatments. *** $p < 0.001$, ** $p < 0.05$, * $p < 0.1$.

We also find that women and older subjects on average report significantly lower production outputs than men and younger subjects, respectively. These effects are also robust when controlling for the other variables (models (2) and (3)). Moreover, when pooling the data of all three treatments we see that women report a production output of zero significantly more often than men (χ^2 -test: $p=0.019$, two-sided). Recall that reporting a zero is a strong indicator of honesty since it results in a payoff of zero. There is an ongoing discussion on whether women tend to lie less than men when payoffs are at stake. Some studies support the notion of the less lying female gender (Ross and Robertson, 2000; Dreber and Johannesson, 2008; Ellingsen et al., 2009; Pruckner and Sausgruber,

2013). However, results from these studies do not always turn out to be significant and there are also studies, which endorse that women lie more than men (Tyler and Feldman, 2004; Tyler et al., 2006) or that there are no differences in lying behavior among genders (Lewis, 1993; DePaulo et al., 1996; Rowatt et al., 1998; Cadsby et al., 2010; Belot and Schröder, 2012). Our study is in line with the first stream of literature. Showing that age is negatively correlated with reported production outputs confirms previous findings on the relationship between age and lying behavior (i.e., see Ross and Robertson, 2000). Given the small variance in age in our student sample ($SD=4.9$) we do not want to overemphasize this result but we consider it quite interesting that the standard result can already be confirmed with our data.

To check for the influence of the Big-Five personality factors, we included four of the five factors into our model. Scale reliability is acceptable for Extraversion (Cronbach's $\alpha=0.69$), Conscientiousness (0.475), Neuroticism (0.587), and Openness (0.559). For Agreeableness scale reliability is unacceptably low (0.107); therefore we refrain from using this factor in our analysis.⁶ Model (3) shows that when controlling for Treatment, Female, and Age, Extraversion and Neuroticism positively predict reported production outputs. Our finding supports previous results that extravert individuals - who are seen as gregarious, assertive, active, self-conscious, and excitement seeking and who like to participate in social interactions - are more inclined to tell lies than introvert persons (Kashy and DePaulo, 1996; Rowatt et al., 1998). There might also be lies born out of insecurity: people who have low self-consciousness and who are anxious may also choose to lie to appear more compliant and more motivated (e.g., Buss and Briggs, 1984; Kashy and DePaulo, 1996) - this inclination might be especially enhanced when those people can disguise their lies. We find support for this conjecture: subjects scoring high in the Neuroticism domain report higher production outputs.

We conclude that team incentives really matter for honest conduct. Individual difference variables are also correlated with subjects' inclination

⁶Rammstedt and John themselves point to a noticeable and substantial loss in reliability and validity in their 2-item Agreeableness scale as compared to larger Big-Five personality measures like, e.g., NEO-PI-R or BFI-44 (see Rammstedt and John, 2007, p.210).

to lie.

3.5 Concluding Remarks

We investigate the influence of two widespread compensation schemes on agents' inclination to lie. By employing a simple experimental design introduced by Fischbacher and Föllmi-Heusi (2013) and controlling for individual difference variables (Female, Age, and Big-Five personality traits) we find that lying is prevalent under both schemes but more pronounced under team incentives than individual piece-rates. This robust finding highlights a so far fairly neglected feature of compensation schemes. It shows that organizations are well advised to be vigilant regarding potentially harming side-effects of compensation schemes, especially when agents can diffuse or displace their responsibility for unethical conduct. Our study shows that besides benefiting others, agents working under team incentives might be particularly prone to lying and deception because it is easier to hide individual wrongdoing under a team than an individual compensation scheme which makes it more difficult to pinpoint a liar in teams. Interestingly, this finding on the corruptive effect of a contingent responsibility shift supports an argument raised by Fischbacher and Föllmi-Heusi (2013). In their seminal paper they show that many subjects do not lie *to the full extent*, i.e., subjects report an outcome of 4 instead of 5 (see also Observation 3). The authors argue that subjects know that honesty might be a favorable trait and if a 4 is assessed differently than a 5 in respect to honesty, it might be reasonable not to lie to the full extent and to try to disguise the lie and appear honest. Similarly, players in our treatment *Team* and players with team compensation in *Team-Mixed* might also see lying to be unethical but they are able to diffuse their responsibility.

3.6 Appendix

Instructions of the Experiment (translated from German)

Statistical Survey

For a statistical survey (independent of the previous experiment) we ask you to answer some brief questions. To determine your payment for answering the upcoming questions, we ask you to roll a die. If you roll a 1, 2, 3, 4, or 5 you receive the diced number as points. If you roll a 6 you receive 0 points. The following applies to you:

{[*Individual*] For answering the questions, the points will be paid to you in Euro.}

{[*Team*] Another subject will be randomly assigned to you. The points of this subject and your points will be added up. For answering the questions, you and the other subject each will be paid half of the sum of these points in Euro.}

{[Team-Mixed, Player with individual compensation] Another subject will be randomly assigned to you. Your points and the points of a random die roll by the computer will be added up. For answering the questions, you will be paid half of the sum of these points in Euro.

The following rules apply for the subject assigned to you: The points of this subject and your points will be added up. For answering the questions, this subject will be paid half of the sum of these points in Euro.}

{[Team-Mixed, Player with team compensation] Another subject will be randomly assigned to you. The points of this subject and your points will be added up. For answering the questions, you will be paid half of the sum of these points in Euro.

The following rules apply for the subject assigned to you: The points of this subject and the points of a random die roll by the computer will be added up. For answering the questions, this subject will be paid half of the sum of these points in Euro. That is, the payment of the other subject does not depend on your die roll.}

Now, please dice your number (the number diced at the first time counts; afterwards you can roll the die several times to check if the die is really fair). Now enter the diced number that you have rolled at the first time:
[]

When you are ready, please fold this sheet once and hold it out of the cubicle. When all subjects are done the sheets will be collected. Afterwards the questionnaire will be handed out to you. We ask you to carefully answer the questions.

{ } = Indicate the phrase that is exclusively employed in the respective treatment, [*Individual*], [*Team*] or [*Team-Mixed*].

4 The Effect of Communication Channels on Lying

4.1 Introduction

Everyday people use a large variety of communication channels to interact with others on and off the job. Think about an employee who has to report the hours he has worked on a project to determine his hourly pay. The actual working time is the employee's private information and can be reported via email or phone. Or imagine a job candidate who has to disclose information to a potential employer, e.g., about his or her current wage level to determine the future income. The information can be requested over an online application system or verbally during an interview. The veracity of private information transmissions might crucially depend on whether the private information is reported verbally, by phone, or via computer. In one situation people may be more likely to lie than in others. This paper investigates how reporting behavior is affected by communication channels. The answers may lead to fruitful insights not only for economics but also for e-commerce, for example when creating online platforms, and organizational designers, for instance when developing corporate reporting structures.

The most natural way of communicating is face-to-face. Communicating parties see and hear each other and can instantly respond. Due to technological improvements during the last decades, people now commonly make use of a large variety of other communication channels. In addition to the classic telephone at home, mobile phones are common. Moreover, text-based and computer-mediated communication, such as e-mails and instant messaging, are established ways of communicating

both at the workplace and in private life. These different communication channels change the communication environment. A fundamental difference is denoted by verbal, such as face-to-face, and nonverbal, for example e-mail, communication. Verbal and synchronous communication is thereby characterized by verbal clues (e.g., tone of voice), visual clues (e.g., body language), and social clues (e.g., status; see Bicchieri and Lev-On, 2011). Non-verbal communication is marked by non-synchronicity and recordability (see Hancock et al., 2004). In line with this differentiation, media richness theory classifies different channels of communication with respect to their richness. A medium becomes richer the more it is able to transmit different types of information from a sender to a receiver (see Daft and Lengel, 1986). Face-to-face communication is considered to be the richest medium as it provides synchronicity and multiple clues systems.

This paper experimentally investigates whether lying is affected by the channel of communication. There are few other studies that look at how communication channels influence lying behavior. Valley et al. (1998) study a bilateral negotiation game with asymmetric information, finding different degrees of trust, truth-telling and efficiency across communication channels. Subjects negotiating face-to-face achieve higher joint benefits due to higher levels of truth-telling than those negotiating by telephone or in writing. However, in their game, subjects had incentives to behave honestly in order to achieve efficient outcomes (see also Valley et al., 2002).

From the field of social psychology, Frank (1988) highlights that physical ‘tell tale clues’, such as facial expressions, blushing and tone of voice, exist under synchronous communication. As these clues are difficult to control for, people may be, for instance, more cautious about lying face-to-face than in writing emails. In line with this, DePaulo et al. (1996) also argue that people will avoid most direct modes of social interaction when telling lies (see, more generally, Vrij (2010), on lying detection).¹

In analyzing people’s diaries, research in communication science comes to findings different than those in social psychology. Hancock et al.

¹In his classic studies on obedience to authority, Milgram (1965) varied the the ‘closeness of authority’: in one condition the experimenter gave instructions face-to-face whereas in the other by phone. In the later condition subjects were significantly less obedient.

(2004) find that synchronicity and recordability of communication primarily matters for the decision of whether to lie or not. They state that the majority of lies occur spontaneously and unplanned. In their analysis, they find that synchronous channels are more prone to deception, while users of recordable and non-synchronous channels seem to be more cautious when lying, as false claims might be reviewable.

Brosig et al. (2003) analyze the effects of different communication channels on cooperation in several standard public good games. The authors vary the communication channel applied in pre-play communication, e.g., auditory or visual channels, either bidirectional or unidirectional. They find that bidirectional face-to-face communication is crucial for enhancing cooperation.² In line with this result, Bicchieri and Lev-On (2007) state that face-to-face communication more easily focuses agents on social norms, i.e., to cooperate in social dilemma games.³

In a coin flip study by Abeler et al. (2014), a representative sample of the German population was called and asked at home to report the outcome of one single coin flip or four coin flips, respectively. Before reporting the outcome of the private coin flip(s), a 20-minute survey-interview was conducted on the phone. In the single coin flip case, participants received a payment if tails comes up. If they were asked to flip the coin four times, each tails flip triggered a payment. The reported outcome and not the actual outcome of the coin flip(s) determined the individual pay. The authors do not find dishonest reporting either in the single or in the four coin flip case. In two laboratory control treatments, subjects reported their outcomes either via phone or computer. While they find significant over-reporting in the laboratory, average reporting between the two communication channels is not significantly different. However, subjects reported the payoff-maximizing outcome, i.e., 4 times tails, more often via computer (see also Fischbacher and Föllmi-Heusi

²Brosig (2002) also underlines the importance of face-to-face communication for cooperation. The author finds that cooperation in a prisoner's dilemma only results when the other player can be identified as cooperative beforehand during a pre-play communication phase. Moreover, she shows that subjects who are classified as cooperative type in a preceding task are better at identifying their partner's willingness to cooperate in the subsequent prisoner's dilemma game.

³Bordia (1997) reviews the effects of face-to-face and computer-mediated communication on group behavior. He finds that groups who communicate via the computer are more prone to uninhibited behavior, e.g., rude and offensive actions.

2013, on ‘partial lying’).⁴

The present experiment employs an experimental design similar to the one used by Abeler et al. (2014), yet it systematically varies four common means of communication, i.e., face-to-face, phone, computer-mediated, and online. Subjects are instructed to privately flip a coin four times. Each time tails comes up, they earn an additional €1 on top of a fixed payment of €7 for completing a required post-experimental questionnaire. After the four coin flips, subjects have to report the number of tails. In comparison to many other studies on lying (e.g., Fischbacher and Föllmi-Heusi, 2013), this experiment is not conducted after another unrelated experiment (see also Houser et al., 2012), but rather as stand-alone session. The compensation structure incentivizes subjects to dishonestly over-report the true number of tails. Across experimental treatments, the applied channel of communication varies between (i) face-to-face, (ii) telephone, (iii) computer-mediated, and (iv) online communication. As the true outcomes of the single coin flips are unknown, aggregate behavior of the reported outcomes is compared to the true distribution of four fair coin flips. Lying is measured by the deviation of the reported outcomes from the expected truthful distribution. The results show that subjects lie independently of the communication channel applied. The extent of lying, however, differs significantly: more subjects lie to the full extent, i.e., they report an outcome of 4 times tails, in treatments with computer-mediated and online communication, whereas partial lying, i.e., reporting an outcome of 3 tails flips, is prevalent in treatments with more personal and synchronous communication. The findings underline the variation in lying costs across decision making environments (see Mazar and Ariely, 2006). Image concerns (e.g., Bénabou and Tirole, 2006; Andreoni and Bernheim, 2009) and social distance arguments (e.g., Bohnet and Frey, 1999) both have possible explanatory value for this variation. Physical clues, which can affect the probability of lies being detected, are also considered as a possible mediating influence (e.g., Frank, 1988).

The next section explains the design of the study in detail. In section

⁴Houser et al. (2012) and Bucciol and Piovesan (2011) conducted laboratory studies with single coin flips and also find dishonest over-reporting. In both studies, however, subjects had to write down the outcome of their private coin flip on a sheet of paper.

3 the theoretical considerations are reviewed and behavioral hypotheses are derived. Section 4 presents the results. Sections 5 and 6 discuss the results and conclude.

4.2 Experiment

4.2.1 Experimental Design

Subjects can earn money by privately flipping a coin four times. Nobody apart from the respective subject, i.e., neither the experimenters nor any other subject, can observe the actual outcome of the four coin flips. Each time the coin lands on tails subjects receive €1. Accordingly, participants can earn €0, 1, 2, 3, or 4 by reporting 0, 1, 2, 3, or 4 tails outcomes. The payoff depends on reported outcomes and subjects can increase their total payoff by lying. After flipping the coin subjects are required to fill in a questionnaire for which they are paid a fixed wage of €7.

The means of communicating the number of tails are systematically varied in four experimental treatments. In the first treatment - *Face-to-Face* (henceforth: *F-t-F*) - subjects report the number of tails to a research assistant in their cabins. In the second treatment - *Phone* - subjects are contacted by the research assistant through a telephone line, i.e., Skype. Each subject in the cabin is equipped with a headset and headphone. Subjects are called by the research assistant and asked to report the number of tails. In the third treatment - *PC-Lab* - no direct verbal communication is employed. Here, subjects are asked to report the number of tails through an entry mask on their computer screen and send it to the research assistant. This treatment resembles e-mail communication practiced daily in organizations. In the fourth treatment - *PC-Online* - subjects participate in an online experiment. Again they are asked to toss a coin four times at home.⁵ Similar to the *PC-Lab* treatment, subjects are asked to report the number of tails over an online entry mask. The treatment variations are summarized in Table 4.1.

Afterwards, subjects answer a comprehensive questionnaire about demographics and personality traits that might help to better understand

⁵The coin-flip task is especially suitable as it can be expected that subjects always have a coin with them.

Table 4.1: Treatment variations

Treatment	Variation of communication channel	n (% female)
<i>Face-to-Face</i>	Reporting outcome face-to-face	60 (50%)
<i>Phone</i>	Reporting outcome via a telephone line	60 (50%)
<i>PC-Lab</i>	Reporting outcome via PC in the lab	60 (50%)
<i>PC-Online</i>	Reporting outcome via PC online outside the lab	66 (45%)

lying behavior.⁶ In addition, subjects are asked what 100 random people would have reported if tails actually had come up zero times (once, twice, three times, four times). Given that filling in the questionnaire might influence subjects' reporting behavior, subjects are asked to flip the coin and report the outcome before filling in the questionnaire.

4.2.2 Experimental Procedures

A total of 246 subjects (with a mean age of 24 and 49% being female) participated in the experiment. The treatments *F-t-F*, *Phone*, and *PC-Lab* were conducted at the 'elfe' laboratory of the University of Duisburg-Essen. The treatment *PC-Online* was conducted online. All subjects were recruited from the same pool of over 2000 students via ORSEE (Greiner, 2003). Treatments *F-t-F*, *Phone* and *PC-Lab* each consisted of 60 subjects, and 66 subjects took part in treatment *PC-Online*. The experiment was computerized using the BoXS software (Seithe, 2012). Each laboratory session involved 12 subjects. In total, 15 laboratory sessions and one online session were conducted in November 2013. Approximately 30% of the subjects were economics or business administration majors. The other 70% were enrolled in different fields of study, such as law and the natural sciences. Subjects were randomly assigned to soundproof cabins. We asked subjects to take a random coin out of their personal wallet to use for the coin flips.⁷ It was also clearly explained that participants would be asked to fill in a 25-minute questionnaire and that they would receive a fixed wage of € 7.⁸ The content of the in-

⁶The 10-item BIG Five inventory (Rammstedt and John, 2007), a survey with respect to the Schwartz' values theory (Schwartz, 1992), and questions on Machiavellianism (Geis and Moon, 1981) were employed.

⁷In case a subject did not possess a coin, we provided a jar of coins from which a personal coin could be taken.

⁸The original instructions are written in German. The instructions provided in the appendix are translated from German into English.

structions was held constant between the treatments, i.e., independent of whether verbal or non-verbal communication was tested, the same wording was employed. In the treatments with verbal communication, i.e., *F-t-F* and *Phone*, the same female research assistant communicated with the subjects following an identical protocol.⁹ In the treatment *PC-Online*, subjects registered for online sessions. At the beginning of the online experiment, participants received an e-mail with a personalized link. Clicking the link routed subjects to online screens identical to the one used in the laboratory sessions of the treatment *PC-Lab*. In the questionnaire, 87% of subjects indicated that they were at home for the *PC-Online* treatment. The other subjects were at other locations, such as in the library or office.

Finally, participants privately received their payoffs. In the treatment *PC-Online*, payments were transferred to subjects' bank accounts. On average, participants earned € 9.77, and the sessions lasted for approximately 35 minutes.

4.3 Theoretical Considerations and Behavioral Hypotheses

Game theory assumes that selfish agents maximize own payoffs and hence misreport their private information if it benefits them (Crawford and Sobel, 1982). Thus, from a standard economic perspective it is not obvious why people should not report four tails flips, independent of the treatment variations. This results in the following (null-)hypotheses.

Hypothesis 1: *Lying to the full extent is observed independent of the channel of communication.*

Research in experimental economics has shown that a large share of subjects misreport private information to their own material advantage (see Croson, 2005, for a review). A considerable share of subjects, however, refrains from lying maximally, which Fischbacher and Föllmi-Heusi

⁹Note that a research assistant was intentionally chosen to communicate with the subjects, as the experimenter might have been biased in his interaction with the subjects. In both treatments the research assistant asked “*How many of your four coin tosses came up with tails?*”.

(2013) label as ‘partial lying’. These findings suggest individual heterogeneity with respect to lying. Theoretical models try to capture these behavioral patterns by considering heterogeneous (psychosocial) moral costs of lying (see, Kartik, 2009; Gibson et al., 2013).

The relevance and the extent of lying costs could be affected by the social environment in which incentives to lie are prevalent (Mazar and Ariely, 2006). The communication channel in place might influence lying costs on at least two levels, though (1) self- and social-image concerns (e.g., Bénabou and Tirole, 2006; Andreoni and Bernheim, 2009) and (2) social distance arguments (e.g., Hoffman et al., 1996; Bohnet and Frey, 1999). Both levels are also connected to the detectability of dishonest behavior, which can be assumed to be different across different channels of communication. Due to physical ‘tell tale clues’, e.g., facial or verbal expressions, detectability is higher under synchronous channels compared to non-synchronous channels (see Frank, 1988; Valley et al., 1998).

First, self- and social-image concerns have been suggested as one possible explanation for partial lying. Some people engage in partial lying in order to maintain a positive image of themselves (e.g., Mazar et al., 2008; Fischbacher and Föllmi-Heusi, 2013). In addition, people want to be perceived as credible and not as greedy, by themselves as well as by others (e.g., Bénabou and Tirole, 2006; Grossman, 2010). In all treatments self-image concerns should be present to the same degree, as subjects have to justify their reported outcome to themselves. The social-image might be influenced by the research assistant’s observation of the subject’s reporting behavior and by the information that is transmitted via the reported outcome. The receiver of the report could form an opinion, i.e., the social-image, about (the honesty of) the sender. It can be assumed that this social-image formation does not depend on the reporting medium. However, the social-image of the subject might be harmed if someone else can more easily detect dishonest behavior due to physical clues observable through synchronous channels of communication. Hence, the interaction of social-image concerns with a higher probability of being detected when lying may increase the relevance of lying costs.

Second, social distance arguments and social norm activation may be suited to explain the effects of different communication channels on lying. Gächter and Fehr (1999) define social distance as the degree of familiar-

ity or anonymity within a social interaction. Hoffman et al. (1996) and Bohnet and Frey (1999) assume that decreasing social distance increases pro-social behavior. The latter authors argue that the identification of the ‘other’ causes more prosociality by assuming that the social distance decreases when “(...) the ‘other’ is no longer some unknown individual from some anonymous crowd but becomes an identifiable victim” (p. 335).¹⁰ With the identification of the other, pro-social norms are activated alongside intrinsic moral norms (Bohnet and Frey, 1999).¹¹ Hence, the closer a social interaction is, the more both motives prevail, i.e., intrinsic moral norms as well as social norms. Under anonymity, however, only the intrinsic motivation to behave morally is decisive.¹² Similar to social-image concerns, social distance is moderated by the detectability due to physical clues. The adherence to social norms may be more pronounced in socially close interactions since norm-deviating actions can be more easily revealed.

Taken together, these lines of reasoning help to derive two alternative hypotheses about behavior in the four treatments.

Hypothesis 2: *More lying is observed under socially more distant channels of communication, i.e., $F(F-t-F) > F(\text{Phone}) > F(\text{PC-Lab}) > F(\text{PC-Online})$, where $F(\cdot)$ denotes the accumulated distribution of the number of lies in a respective treatment.*

Moreover and related to findings from other experiments (e.g., Fischbacher and Föllmi-Heusi, 2013; Abeler et al., 2014), it can also be hypothesized that more partial lying, i.e., reporting the outcome 3, than lying to the full extent will be observed in treatments *F-t-F* and *Phone* compared to treatments *PC-Lab* and *PC-Online*.

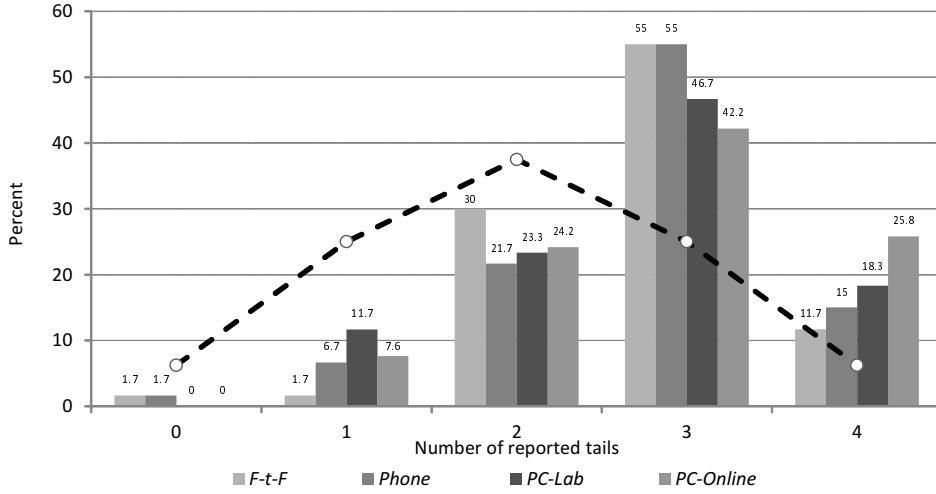
Hypothesis 3: *Partial lying is observed more frequently under channels of communication that are socially more close.*

¹⁰Charness and Gneezy (2008) specify social distance as the extent of physical and emotional proximity induced by a situation (see also Charness et al., 2007). They state that the influence of social distance is especially relevant for e-commerce.

¹¹By following Bicchieri (2006), social norms can be defined as rules or standards of behavior in a reference network such that individuals prefer to conform to them.

¹²Schram and Charness (2011) argue that an increasing number of decisions are being made privately behind a computer screen, which underlines the importance of intrinsic moral norms.

Figure 4.1: Frequencies of reported tails



Note: The dashed line corresponds to the expected truthful distribution.

4.4 Results

4.4.1 Behavior of Subjects

Figure 4.1 depicts the distribution of reported outcomes across treatments. The dashed line represents the expected frequency if every participant reported the true outcome of the four coin flips. The main results are also summarized in Table 4.2.

***Observation 1:** In all treatments, the overall distribution of reported outcomes is significantly shifted to the right of the truthful distribution.*

Figure 4.1 reveals that in all treatments the observed distributions of reported outcomes differ from the theoretical benchmark. The most frequent outcome reported is 3, i.e., a majority of subjects refrains from reporting the payoff-maximizing outcome of 4. The distribution of reported outcomes in all four treatments is significantly different from the truthful distribution (Kolmogorov-Smirnov tests, all four $p < 0.01$). This finding is also confirmed by binomial tests. In the treatments *Phone*, *PC-Lab* and *PC-Online*, the frequencies of reporting 1 or 2 are significantly below the expected distribution, whereas the frequencies of reported outcomes of 3 and 4 are significantly above the honest distribution (see Table

Table 4.2: Overview of results

Treatment	n	\bar{r}_i	Belief \bar{r}_i	Reported outcome r_i (rel. frequencies)				
				0	1	2	3	4
<i>F-t-F</i>	60	2.73	2.85	0.02	0.02 ⁻⁻⁻	0.3	0.55 ⁺⁺⁺	0.12
<i>Phone</i>	60	2.75	2.89	0.02	0.07 ⁻⁻	0.22 ⁻⁻⁻	0.55 ⁺⁺⁺	0.15 ⁺⁺
<i>PC-Lab</i>	60	2.71	2.80	0.00 ⁻	0.12 ⁻⁻	0.23 ⁻⁻	0.47 ⁺⁺⁺	0.18 ⁺⁺⁺
<i>PC-Online</i>	66	2.86	2.86	0.00 ⁻⁻	0.08 ⁻⁻⁻	0.24 ⁻⁻	0.42 ⁺⁺	0.26 ⁺⁺⁺
Honest distribution		2		0.06	0.25	0.38	0.25	0.06
JT test				$p=0.097$	$p=0.065$	$p=0.729$	$p=0.051$	$p=0.017$

Notes: n stands for the number of observations. \bar{r}_i is the average reported output. Plus and minus signs display the significance of a one-sided binomial test indicating that the observed relative frequency is smaller (larger) than the expected frequency ⁻(⁺)=10%-level, ⁻⁻(⁺⁺) =5%-level, ⁻⁻⁻(⁺⁺⁺)=1%-level. JT test stands for applying a Jonckheere-Terpstra test for ordered alternatives (p -values refer to one-sided tests).

4.2 for details). In the *F-t-F* treatment, the reported frequencies for outcomes 1 and 3 are also statistically distinguishable from the expected frequencies. In this treatment, however, the frequencies of reporting 0, 2 and 4 are statistically indistinguishable from the truthful distribution. The average reported outcomes are not statistically different between the treatments.

Observation 2: *The fraction of subjects reporting the outcome of 4 increases the more distant the means of communication is.*

Focusing only on the reported outcome of 4, Figure 1 displays differences in reporting this outcome across treatments. According to a Jonckheere-Terpstra test for ordered alternatives, reporting the highest possible outcome is more likely as the anonymity of the communication medium increases ($p=0.017$, one-sided).

Pairwise comparisons of the fraction of subjects reporting the outcome 4 yields a significant difference between treatments *F-t-F* and *PC-Online*, i.e., 12% in *F-t-F* and 26% in *PC-Online* ($p=0.04$, χ^2 -test, two-sided). No statistical differences exist for other pairwise treatment comparisons (see also Table A4.1 in the appendix for probit regressions on reporting the outcome 4).

In treatment *PC-Online*, 30% of the women report the outcome 4,

whereas in the other treatments only 10% or less of the women report the highest possible outcome (see also Figure A4.1 and A4.2 in the appendix). Pairwise χ^2 -tests of the fractions of women reporting the outcome 4 yield significant differences between treatment *PC-Online* and the other treatments ($p=0.053$ between *PC-Online* and *F-t-F* as well as *Phone*, and $p=0.02$ between *PC-Online* and *PC-Lab*, two-sided). Comparisons between reporting behavior of men and women reveal no statistical differences.¹³

Observation 3: *Partial lying is more pronounced in treatments where the means of communication is less anonymous.*

Subjects across treatments most frequently report outcome 3. According to a Jonckheere-Terpstra-test, reporting this outcome is more likely as the anonymity of the communication channel decreases ($p=0.051$, one-sided). Pairwise treatment comparisons show no statistical differences.

Observation 4: *Being high in Conformity values has a negative effect on reporting high outcomes.*

Ordered logistic regressions of the reported outcomes with several covariates were run that are displayed in Table 4.3. Model (1), dummy variables of the treatments were included, which have no significant effect on reported outcomes. Moreover, neither age nor gender have a significant effect on reporting outcomes (model 2). Model (3) includes a dummy for Conformity values from the value theory of Schwartz (1992). Subjects who score higher in Conformity report significantly smaller outcomes.¹⁴ According to the definition of the conformity trait, people who score highly in Conformity values avoid actions and intentions that could upset or harm others and violate social norms and expectations (Schwartz, 1992). Other characteristics like religiousness, income, risk attitudes and Machiavellianism are not significantly associated with the reported outcome (see Table A4.2 in the appendix).

¹³In contrast, Dreber and Johannesson (2008), for example, find that men lie more often compared to women.

¹⁴According to Schwartz's model (1992), people distinguish between at least ten basic values. These values are Universalism, Benevolence, Tradition, Conformity, Security, Power, Achievement, Hedonism, Stimulation, and Self-Direction (Schwartz, 1992).

Table 4.3: Regression analysis of reported outcomes

Dependent variable:					
Number of reported tails	(1)	(2)	(3)	(4)	(5)
<i>Phone</i>	0.102 (0.303)	0.072 (0.307)	0.057 (0.312)	0.103 (0.345)	0.037 (0.349)
<i>PC-Lab</i>	0.014 (0.330)	-0.006 (0.332)	-0.001 (0.349)	0.020 (0.349)	-0.015 (0.354)
<i>PC-Online</i>	0.335 (0.333)	0.337 (0.337)	0.232 (0.341)	0.108 (0.347)	0.245 (0.359)
Age		-0.034 (0.031)			-0.041 (0.033)
Gender (1 if female)		-0.083 (0.243)			-0.074 (0.252)
Conformity			-1.207** (0.415)		-1.245** (0.485)
Belief about other subjects				4.927*** (0.757)	4.676*** (0.766)
Observations	246	246	246	242	242
Pseudo <i>R</i> -squared	0.002	0.004	0.014	0.084	0.087

Notes: Ordered logit estimates. Robust standard errors in parentheses. Reference group is treatment *F-t-F*. The number of observations differs because of missing values. *Phone*, *PC-Lab*, and *PC-Online* represent dummy variables for the respective treatments. "Belief about other subjects" is the belief about the share of other subjects who report to have flipped more tails than they actually did and who then report the outcome 4 tails (see section 4.4.2 for details). Significances at the 1, 5, and 10 percent level are denoted by ***, **, and *, respectively.

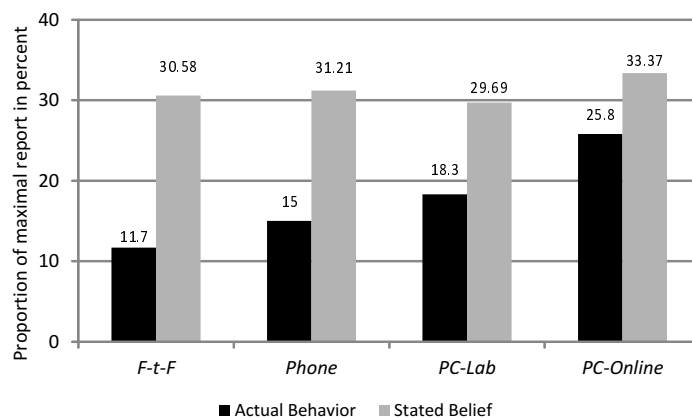
4.4.2 Beliefs about Behavior of Other Subjects

In the questionnaire subjects state what they believe other participants would have reported. Specifically, subjects are asked: "Imagine what 100 other participants would have reported if they had actually flipped tails zero (*one, two, three, four*) times." For every possible actual outcome they allocated 100 fictitious participants to the five possible reported outcomes (i.e., reporting an outcome of 0, 1, 2, 3, or 4 times tails). Thus in total, subjects filled in 25 belief questions, five for each possible outcome (see appendix for questions on beliefs). The belief elicitation was not incentivized.

On average, 53.2% of the subjects state that they believe that other participants report higher outcomes than the actual honest outcomes. Table 4.2 includes beliefs about average reported outcomes.¹⁵ Across

¹⁵Beliefs about average reported outcomes are calculated as follows: For each respective *actual* outcome of the four coin flips, an average individual belief about the

Figure 4.2: Actual behavior and beliefs on payoff-maximizing reports



Note: See also Figure A4.3 in the appendix for the full distribution of beliefs over the reported outcomes.

treatments, average beliefs are not statistically different from the actual reported outcomes.

Figure 4.2 focuses on beliefs about the share of participants who report the payoff-maximizing outcome, i.e., reporting the outcome 4. The beliefs about reported outcomes (gray bars) are compared with actual proportions (dark bars). In all four treatments subjects believe that others report the outcome 4 more often than they actually do. In contrast to actual reporting of outcome 4, no differences in stated beliefs can be observed between the treatments.

In Table 4.3 model (4), the reported outcome of the four coin flips is regressed on treatment dummies and beliefs about the share of other subjects who dishonestly reported the outcome 4. Subjects who believe that others dishonestly report the highest outcome also report higher outcomes themselves.¹⁶

corresponding *reported* outcome is calculated which is weighted with the respective theoretical probability of occurrence. The sum of these five values determines the belief about the average total reported outcome shown in Table 4.2.

¹⁶The meaning of ex-post stated beliefs have to be examined cautiously. Subjects might justify their (dishonest) behavior with a stated belief that others also exaggerated their reports (e.g., Abeler et al., 2014). The other direction might be also true, i.e., high beliefs about others' reports may induce subjects to also report higher outcomes (e.g., López-Pérez, 2012).

4.5 Discussion: Partial Lying and Communication Channels

Fischbacher and Föllmi-Heusi (2013) find that a majority of participants do not overstate the outcome of a private die roll to full payoff-maximizing extent. This finding is replicated by other studies (e.g., Conrads et al., 2013) and often explained with the idea of self-concept maintenance, i.e., people lie partially because it does not require changing one's self-image as an honest person (Bénabou and Tirole, 2002; Mazar et al., 2008). In the coin flip task employed in this study, the reported outcome of 3 can be interpreted a partial lie in case the actual outcome is below 3.

As mentioned by Fischbacher and Föllmi-Heusi (2013), people may want to be perceived as credible by themselves and by others. Across all treatments, subjects are able to maintain a positive self-concept by lying partially. However, higher identification of subjects through a specific communication channel may lead to greater activation of the social norm not to lie to the full extent (see also Biel and Thøgersen, 2007). This activation does not prevent subjects from lying, but rather only seems to restrain them from payoff-maximizing lying.¹⁷ In the treatment *PC-Online*, subjects act under the highest degree of anonymity, whereby only intrinsic moral norms - and not social norms - may be active (see Schram and Charness, 2011). Here, less partial over-reporting and more over-reporting to the full extent is found.

The beliefs about other people's behavior, however, do not differ across treatments. Subjects believe that others over-report the outcome of the four coin flips, which could be interpreted as the social norm in the given experimental setup. As pointed out by Bicchieri (2006), these *empirical expectations*, i.e., the beliefs about the behavior of others, may not be sufficient to motivate compliance with this social norm. The author further states that conformity to a social norm also depends on *normative expectations*, i.e., the beliefs about what other people believe one should do. Independent of the treatment variations, it can be assumed that the normative expectation is not to lie at all. However, as shown by Bicchieri and Xiao (2009), empirical expectations are usually an accurate

¹⁷Also Abeler et al. (2014) found significantly less lying to the full extent in their phone treatment in the laboratory than in the treatment with computer-mediated communication.

predictor of actual behavior. This reasoning is best reflected in treatment *PC-Online*, where actual behavior is close to the empirical expectations. In the other treatments, the channel of communication might prevent subjects from completely following their empirical expectations since less payoff-maximizing over-reporting is observable.

More broadly, an important reason for the observed behavioral pattern might be that subjects in treatments with synchronous interaction, i.e., in *F-t-F* and *Phone*, are aware of their physical tell tale clues (see Frank, 1988), i.e., their cadence of voice and facial expressions which might reveal dishonest behavior. Although subjects in these treatments are probably still lying, they might feel to exercise better control over their tell tale clues when lying partially. By contrast, under computer-mediated communication they may not fear being detected through physical reactions.

4.6 Conclusion

The influence of four common communication channels on subjects' inclination to lie is investigated by systematically varying the means by which private information is reported. In a simple experimental coin flip game with incentives to overstate private information, lying seems to be prevalent under all channels of communication. Subjects do not differentiate between lying face-to-face, on the phone, or via the computer. Lying to the full extent is, however, most pronounced in online computer-mediated communication when subjects are at home. These findings underline that the relevance and the extent of lying costs seem to differ across different environmental settings (see Mazar and Ariely, 2006), which is a rather neglected contextual feature in the economics literature. On a more general level, the findings are in line with other studies showing that more pro-social behavior can be observed under more direct channels of communication (e.g., Valley et al., 1998; Brosig et al., 2003). Nevertheless, the results do not allow the identification of a conclusive reason for the observed behavior.

Social distance arguments and social norm activation could explain behavior, as the distance varies across communication channels and less over-reporting to the full extent is observed under more socially-close

channels (Bohnet and Frey, 1999; Charness and Gneezy, 2008). By reducing social distance, both the compliance to one's own moral norms *and* social norms may lead to less lying to the full extent compared to situations with larger social distance, where only intrinsic moral motives are prevalent.¹⁸ This explanation is also influenced by the probability of being lying detected, which differs across communication channels due to physical reactions, for example, cadence of voice (Frank, 1988). Socially distant interactions are characterized by a smaller probability of being detected since no physical clues can be observed.

The prevalence of electronic and impersonal communication channels is unstoppable. Internet businesses and virtual markets are growing. Designers of e-commerce platforms and organizational reporting structures should, however, be aware of possible drawbacks of specific channels of communication. Given the findings of this study, employers, for example, should ask job candidates face-to-face about their current income, as this statement might be overstated to the smallest extent in comparison to online application systems.

A fruitful avenue for future research may be to more deeply analyze how social norms and pro-social behavior are activated through specific contextual conditions. More generally, it seems to be important to better understand the effects of depersonalization on social and business interactions.

¹⁸The social distance might even be further reduced the longer individuals interact with each other since the familiarity increases. For instance, in the study of Abeler et al. (2014) subjects in one treatment talk to an interviewer for 20 minutes before reporting the outcome of the coin toss.

4.7 Appendix

Table A4.1: Probit regressions on reporting the outcome 4

Dependent Variable:	(1)	(2)
Reporting four tails flips	Mrg. Effects	Mrg. Effects
<i>[F-t-F]</i>	[0.117***]	[0.198]
	[0.584]	[0.708]
<i>Phone</i>	0.041	0.036
	(0.079)	(0.079)
<i>PC-Lab</i>	0.079	0.071
	(0.082)	(0.080)
<i>PC-Online</i>	0.154**	0.159**
	(0.083)	(0.086)
Age		-0.005
		(0.007)
Gender (1 if female)		-0.073
		(0.048)
Observations	246	246
Pseudo <i>R</i> -squared	0.020	0.030


Notes: Coefficients display marginal effects. Reference group (constant) is treatment *F-t-F* (reported in square brackets). *Phone*, *PC-Lab*, and *PC-Online* represent dummy variables for the respective treatments. Significances at the 1, 5, and 10 percent level are denoted by ***, **, and *, respectively.

Table A4.2: Additional regression analysis of reported outcomes

Dependent variable: Number of reported tails	(1)	(2)	(3)	(4)	(5)	(6)
<i>Phone</i>	0.102 (0.303)	0.072 (0.307)	0.057 (0.312)	0.103 (0.345)	0.037 (0.349)	0.059 (0.353)
<i>PC-Lab</i>	0.014 (0.330)	-0.006 (0.332)	-0.001 (0.349)	0.020 (0.349)	-0.015 (0.354)	-0.029 (0.355)
<i>PC-Online</i>	0.335 (0.333)	0.337 (0.337)	0.232 (0.341)	0.108 (0.347)	0.245 (0.359)	0.252 (0.369)
Age		-0.034 (0.031)			-0.041 (0.033)	-0.0529 (0.035)
Gender (1 if female)		-0.083 (0.243)			-0.074 (0.252)	-0.045 (0.276)
Conformity			-1.207** (0.415)		-1.245** (0.485)	-1.322*** (0.499)
Belief about other subjects				4.927*** (0.757)	4.676*** (0.766)	4.797*** (0.797)
Religiousness						0.019 (0.076)
Income						0.161 (0.159)
Risk						0.008 (0.101)
Machiavellism (1 if above median)						0.027 (0.269)
Observations	246	246	246	242	242	242
Pseudo <i>R</i> -squared	0.002	0.004	0.014	0.084	0.087	0.084


Notes: Ordered logit estimates. Robust standard errors in parentheses. Reference group is treatment *F-I-F*. The number of observations differs because of missing values. *Phone*, *PC-Lab*, and *PC-Online* represent dummy variables for the respective treatments. "Belief about other subjects" is the belief about the share of other subjects who report to have flipped more tails than they actually did and who then report the outcome 4 tails (see section 5.2 for details). For risk attitudes general risk questions of the GSOEP were included. Values of Machiavellism are included through a dummy variable for subjects who score above average in Conformity. Significances at the 1, 5, and 10 percent level is denoted by ***, **, and *, respectively.

Figure A4.1: Frequencies of reported tails by men




Note: The dashed line corresponds to the expected truthful distribution.

Figure A4.2: Frequencies of reported tails by women



Note: The dashed line corresponds to the expected truthful distribution.

Figure A4.3: Beliefs about frequencies of reported tails



Note: The dashed line corresponds to the expected truthful distribution. See also Figure 4.2 in the section 4.4.2.

Instructions of the Experiment (translated from German)

Welcome to today's experiment and thank you for participating. Now, please carefully read the instructions. In case of questions arise, do not hesitate to raise your hand. We can personally assist you.

In the following, we ask you to fill in a questionnaire. Your information will be anonymized. It will take 20-25 minutes to fill in the questionnaire. For filling it in you will receive a fixed payment of €7.

Before you start filling in the questionnaire, we ask you to conduct four coin tosses. Therefore, please take a coin out of your wallet. {[*Face-to-Face*, *Phone*, *PC-Lab*: In case you do not have a coin with you, we will provide you with a coin.]} On one side of the coin there is tails and on the other side heads. Every time tails comes up, you receive one additional Euro. You receive €0 if tails never comes up; you receive €1 if tails comes up once; you receive €2 if tails comes up twice; you receive €3 if tails comes up three times; and you receive €4 if tails comes up four times.

The payoff from the four coin tosses will be added to your fixed payment for filling in the questionnaire. Please now toss your coin four times (the first four tosses count).

{*Face-to-Face*: Soon, we will come to your cabin. We will ask you to tell us the number of coin tosses that came up with tails.}

{*Phone*: Soon, we will contact you in your cabin via Skype. Therefore, please put on your head-set. We will ask you to tell us the number of coin tosses that came up with tails.}

{*PC-Lab* & *PC-Online*: Soon, we will be connected via an online-connection. We will ask you to tell us the number of coin tosses that came up with tails.}

After you told us the number of coin tosses that came up with tails, the questionnaire will be shown on your computer screen. Please then fill in the questionnaire carefully.

{ } = Indicate the phrase that is exclusively employed in the respective treatment, [*Face-to-Face*], [*Phone*], [*PC-Lab*], or [*PC-Online*].

Questions for Belief Elicitation in the Questionnaire

These questions relate to the potential behavior of other participants in the experiment who find themselves in the same situation you were in.

What do you think: How many times would 100 randomly selected people, who actually tossed tails zero [*one, two, three, four*] times, report tails? Therefore, please allocate these 100 people to the following five possible answers:

report tails 0 times.

report tails 1 times.

report tails 2 times.

report tails 3 times.

report tails 4 times.

5 Fraud at the Kiosk - Field Evidence from Buying Candy

5.1 Introduction

The question of how markets affect individual morality has been debated ever since economics emerge as a scientific field (Smith, 2010; Sandel, 2012, 2013, and Falk and Szech, 2013, for recent experimental findings). What has not been thoroughly debated is the fact that certain market characteristics inflate the risk of moral failures. One example of this can be found in “credence goods” markets, in which informational asymmetries exist between sellers and buyers. In such credence goods markets, customers typically face the risk of being either overcharged or overtreated. Credence goods have most prominently been discussed in the realms of automobile repair or medical services, where customers usually suffer from inferior knowledge compared to expert sellers. As customers simply do not know what they need nor cannot find out due to their lack of expertise, expert sellers can either charge them more than is justified or sell them more than is needed. More recently, the discussion was extended towards professional services, such as computer programming, where the risks seem severe due to rather extreme knowledge differences and high costs of changing the supplier. Darby and Karni (1973) were the first to use the term credence good. They systematically explored the behavior of customers and sellers in credence goods markets and examined possible market equilibria. Subsequent theory (e.g., Dulleck and Kerschbamer, 2006) and experiments (e.g., Dulleck et al., 2011; Balafoutas et al., 2013) highlight the particular risks for customers, which typically involve two above mentioned distinctive threats: overcharging

and overtreatment (see also Domenighetti et al., 1993; Wolinsky, 1993; Schneider, 2012).

The present research addresses an important gap in the literature regarding fraudulent behavior in such credence goods markets. While typically behavior in these markets cannot clearly be attributed to overcharging or overtreatment, we examine a product that - by strict definition - does not qualify as a credence good, but shares essential characteristics that permit overcharging while precluding overtreatment. Additionally, we are able to address the stability of seller behavior across customer encounters and analyze discrimination against particular customers (which we vary in terms of status and purchase size) in two field experiments. The motivation is closely tied to recent experimental evidence showing that professionals in credence goods industries are much more prone to overtreat compared to naive student samples in a laboratory experiment. Beck et al. (2014) find that – possibly due to decision heuristics – car mechanics in a laboratory setting systematically overtreat customers more than students. Therefore, they suggest that overtreatment is "learned", possibly as a result of certain institutional rules (e.g., high warranties or high legal risks of undertreatment). To complement this research, we investigate if certain sellers still systematically overcharge once controlled for overtreatment. We therefore rely on a (non credence good) market that rules out overtreatment and reduces the risk to overcharging faced by customers, while maintaining a setting that one could label "quasi-credence good". To achieve this, we exploit the existence of candy-by-weight pricing schemes frequently used in kiosks in Cologne (Germany) as it allows a panel to be set up in the field to address seller behavior in a repeated setting. Importantly, loose candy is often weighed on an unobservable scale by sellers. If at all, buyers can only check on a personal scale at home directly after the purchase whether they have been charged the correct amount, which gives candy sales a credence good attribute, as it invites overcharging due to informational asymmetry regarding the true weight of the candy.

In this setting, we are able to address two important aspects relevant to credence goods markets. First, we investigate the generalizability of "standard" results, namely whether customers are overcharged at all. To do this, we contrast the behavior of sellers relying on hidden scales to

sellers using openly visible scales (an alternative used in some stores; Experiment 1). Subsequently, making repeated purchases within one hour in the same kiosk with a hidden scale (Experiment 2), we evaluate how robustly sellers overcharge their customers, thereby addressing an important aspect of such markets that - to our knowledge - has not yet been analyzed. We rely on a field experiment to boost external validity of our results and to measure subjects behavior in a naturally occurring environment without their awareness of being monitored (see Beck et al., 2014, and List, 2006, for the general advantages of field experiments).

We generally replicate previous results (Experiment 1) and – importantly – find support for the conclusion that a characterization of sellers is possible, which further indicates some level of behavioral consistency by sellers across purchases (Experiment 2), as well as their discrimination against particular high status buyers that purchase large quantities. In total, our field experiments yield three core findings: First, consistent with the literature, overcharging only occurs in kiosk with hidden scales where customers can be secretly overcharged. Thus, we find that overcharging occurs once overtreatment is ruled out and therefore other explanations have limited applicability (i.e., fear of warranties or best-intentions). In these kiosks, customers were overcharged in 40 (Experiment 1) to 44 (Experiment 2) percent, compared to below 10 percent in kiosks using scales visible to the public. Second, the occurrence of overcharging partly depends on contextual variations, e.g., the amount of candy purchased or the physical appearance of the customer as a representative of a high status group (see theories on distributional preference from Fehr and Schmidt, 1999, and Bolton and Ockenfels, 2000). Third, the data suggest that many sellers could be classified as either honest or cheaters. We find that sellers to a certain degree behave consistently in the repeated purchases, i.e., we were able to classify a majority of sellers as members of a particular category, e.g., being either consistently accurate or fraudulent.

5.2 A Brief Description of Kiosks

As many readers may be unfamiliar with the nature of the kiosks used in the experiments, we briefly describe the market structure of such kiosks.

Quite common in Cologne’s urban landscape, kiosks (“Büdchen” in the local language) are small-sized shops for newspapers, magazines, beverages, and a small range of groceries. Their existence stems from the fact that grocery stores’ opening hours used to be heavily regulated in Germany. For people interested in buying newspapers, cigarettes, or drinks outside of general service hours (often 8 a.m. - 6:30 p.m.), kiosks were the only option.

Despite some deregulation of opening hours in the past decade, kiosks have survived and are a vital part of the city’s shopping scene. Typical markets are smaller than 20 square meters and are often owner- or family-operated, many of them by first or second generation immigrants to Germany. For the present experiments, we relied on markets around the city center and some popular residential areas directly adjacent to the city center (see Figure A5.1 in the appendix for a map). As our good of interest, we use assorted candy that is offered in a plethora of stores. Paying for candy follows one of two general pricing schemes. Either candy is counted and paid for by the item (e.g., 5 cents for small pieces, 10 cents for larger pieces) or candy is weighed and charged per 100 grams (i.e., most markets charge €1 per 100 grams). We focus our study on the latter pricing scheme.

When buying candy in this type of kiosks, customers can choose between several pieces of candy and put them in a little paper bag on their own. Once they finish their candy selection, customers hand over the bag to the seller, who weighs the candy to determine the price. In Experiment 1, we exploit the fact that in some markets the scale is clearly visible from the customer’s perspective, while in some it is not.

One important characteristic of our setting is that it does not allow for best intentions as an excuse for fraud (e.g., avoidance of warranties or overtreatment with best intentions). As discussed by Balafoutas et al. (2013), generous overtreatment with best intentions can be prematurely interpreted as fraud. For example, a taxi driver might be accused of illegitimate detours while in fact having best intentions to take a longer, but faster route. Similarly, a car mechanic (or physician) might be overly cautious by conducting unnecessary actions with the best intention of prolonging the lifetime of a car (patient) or ruling out the risk that a problem occurs again soon. Therefore, the candy-by-weight pricing scheme pro-

vides a natural setting to address the question of overcharging in the field while ruling out overtreatment due to best intentions, a frequent confounding factor of fraudulent behavior in credence goods markets.

5.3 Experiment 1: Overcharging When Selling Candy

Experiment 1 was designed to test the basic assertion that candy sales are prone to overcharging. It employs a quasi-experimental field setup in the city of Cologne. The peculiarity we exploit in Experiment 1 is the position of the scale in kiosks used to weigh candy. In 37 of 82 visited markets, the location of the scale is publicly visible and the outcome of the weighing is observable to customers.¹ Henceforth, this situation is denoted as *Public Scale*. In the remaining 45 markets, the outcome of the scale cannot be observed by the buyer. In this case, the buyer cannot (immediately) verify whether the calculated amount matches the real weight or not. Hence, the seller can easily overcharge the buyer.² In the following, this situation is denoted as *Hidden Scale*. The logical hypothesis is that the risk of being overcharged is higher in markets where the sellers of candy weighs with hidden scales compared to those weighing with public scales.

5.3.1 Procedure

In each of 82 markets, the same male buyer entered the market and collected about 150 grams of candy (i.e., roughly 15 similar pieces). The bag was then handed over to the seller, who weighed it on a digital scale to determine the price. The price was paid by the buyer, after which he left the kiosk. Interaction and appearance was held neutral and con-

¹With the help of research assistants, we scanned the entire city center and adjacent residential areas for kiosks using the pay-per-weight pricing scheme. The sample consists of all identified markets that use this pricing scheme as of November 2011. After this initial scan we predetermined routes of visits that were used to gather data for the experiment. All purchases followed predetermined routes and protocols to minimize unobservable confounding factors.

²Note that overtreatment and undertreatment is also not possible in our setting since the buyer chooses the candy by himself.

stant during all purchases.³ The purchases followed a predefined route, and after leaving the buyer marked the bag with a sticker that uniquely identified the market and the price. Furthermore, 30 minutes after the purchase, two student assistants conducted a questionnaire asking sellers for demographics and personality traits (see Rammstedt and John, 2007). The relationship between the candy sales and the questionnaire was not disclosed to sellers.⁴ After the purchases, the candy-bags were weighed again and the difference between price paid and the correct price was calculated to render the key variable of the experiment. All purchases took place on weekdays (Tuesdays-Thursdays) evenings of November 2011.

5.3.2 Results and Discussion

In the following, the difference between the price paid and the actual price is denoted as Δ_i which can either be positive, i.e., the customer is overcharged, or negative, i.e., the customer is undercharged ($i = 1, \dots, 82$). In case Δ_i takes the value of zero the charged price is correct. Consistent with our initial hypothesis, we find our product to be more expensive in *Hidden Scale*. On average, the charged price is 1.69 cents above the correct price in *Hidden Scale* and -4 cents below the correct price in *Public Scale*. This difference turns out to be significant ($p=0.019$; Fisher-Pitman permutation test, one-sided). The average size of overcharging (conditional on overcharging) is 12 cents in *Hidden Scale* and 5 cents in *Public Scale*. The average size of undercharging (conditional on undercharging) is 9 cents in markets using a public scale and 6 cents in markets employing a hidden scale (see also Tables A5.2 and A5.3 in the appendix for more detailed descriptive statistics).

In the following, the occurrence of honesty, overcharging, and undercharging is analyzed. Table 5.1 summarizes the results. There are differences between markets in the probability of being overcharged. In *Public Scale*, 3 out of 37 sellers (8%) overcharge the customer. In contrast, in *Hidden Scale*, 18 out of 45 sellers (40%) overcharge. This difference

³In fact, the customer just said “hello” when entering the kiosk, “that’s it” when giving the candy bag to the seller, and “goodbye” when leaving.

⁴The student assistants asked the seller to participate in a survey that aims to assess the situation of kiosks in Cologne. 10 of 37 sellers (27%) using open scales and 27 of the 45 (60%) sellers employing hidden scales filled in the survey. See Table A5.1 in the appendix for the results of questionnaire.

Table 5.1: Overview of results from Experiment 1

Treatment	n	Δ_i (means)	% honest	% overcharging	% undercharging
<i>Hidden Scale</i>	45	1.69 √**	9% ∧+++	40% √+++	51% ∨
<i>Public Scale</i>	37	-4	43%	8%	49%

Notes: Δ_i denotes the difference between the price paid and the actual price. Stars display the significance of a one-sided Fisher-Pitman permutation test (*=10 %-level, **=5 %-level, ***=1 %-level). Plus and minus signs display the significance of a one-sided χ^2 -test (-(+)=10 %-level, --(++)=5 %-level, ---(+++)=1 %-level).

is highly significant ($p < 0.01$, χ^2 -tests, one-sided). Thus, we find more sellers to be honest in kiosks employing public scales.

Next, we analyze the frequency of undercharging. In kiosks employing a *Hidden Scale*, 23 out of 45 sellers (51%) are generous. In kiosks using a *Public Scale*, 18 out of 37 sellers (49%) undercharge their customers. This difference is not statistically significant.

Finally, 37 of the 82 sellers filled in the post-experimental questionnaire (27 in *Hidden Scale* and 10 sellers in *Public Scale*). Comparing the data of sellers in kiosks with hidden scales with sellers with public scales, we do not find significant differences with respect to demographics and personality traits (see Table A5.1 in the appendix).⁵ Interestingly, the price charged by sellers who fill in the questionnaire is on average 2.05 cents above the correct price and -3.28 cents below the correct price in case sellers denied filling it in ($p = 0.029$, Fisher-Pitman permutation test, one-sided).⁶

Summarizing, we find significant differences in overcharging between kiosks with *Hidden Scale* and *Public Scale*. By contrast, there were no observable differences in undercharging. Our results replicate previous experimental evidence that some degree of dishonest behavior (e.g., fraudulent overcharging) can occur in the candy market. Therefore, Experiment 1 contributes to the generalizability of the “typical” result, which highlights the risk of being overcharged in credence goods markets. Moreover,

⁵In Table A5.4 in the appendix, Δ is regressed on seller demographics and character traits leading to no significant effects.

⁶A possible explanation for this effect could be that sellers who defrauded their customers when selling candy clean their bad conscience by filling the questionnaire (e.g., see Monin and Miller, 2001, on moral licensing, and Gneezy et al., forthcoming, on conscience accounting).

findings in Experiment 1 allow us to use the candy-paradigm to examine behavioral consistency and discrimination against particular buyers in Experiment 2.

5.4 Experiment 2: Behavioral Consistency of Sellers and Discrimination of Buyers

While Experiment 1 was designed to test the basic assertion that kiosks selling candy can qualify as an example of overcharging, Experiment 2 was designed to test the main research question regarding the behavioral consistency of sellers in such kiosks, along with potential discrimination against particular buyers. Therefore, we designed an experiment informed by previous research (Ebeling et al., 2012; Balafoutas et al., 2013) and counterbalanced two aspects in a total of four consecutive purchases in kiosks with hidden scales. First, we manipulated the appearance of the buyer, making him appear either *rich* or *poor*. The seller’s perception of the customer’s income was implemented by different apparel. Customers (all males) intended to be perceived as having a high income and were dressed in a suit with a tie and were equipped with a briefcase, while low-income customers were dressed shabbily and unshaved and carried a backpack. We implemented this manipulation to test whether seller behavior might be driven by distributional preferences (e.g., Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Kerschbamer et al., 2009). Additionally, we varied the amount of candy bought, either a *large* amount of candy, i.e., 150 grams, or *small* amount, i.e., 50 grams, to vary the probability by which fraud could be detected by the customer.⁷ Hence, in Experiment 2 we implemented a 2x2 experimental design with four treatments, i.e., *Rich-Large*, *Rich-Small*, *Poor-Large*, and *Poor-Small*. Table 5.2 summarizes the experimental treatments of

⁷Whether this was indeed the case was tested in an independent laboratory experiment. We asked 64 participants to give us their weight-estimates of a 50 grams ($n=32$) vs. a 150 grams ($n=32$) bag of candy. Estimations were incentive-compatible. The results suggest that estimates of the large bag are more inaccurate than estimates of a small bag, i.e., standard deviations are larger estimating the larger bag ($SD=106.18$ grams) compared to the small bag ($SD=29.39$ grams; $p<0.01$, variance ratio test). Therefore, it can be assumed that it is less risky for sellers to add a price-premium on a large compared to small a bag of candy. See also Table A5.5 in the appendix.

Table 5.2: Treatments of Experiment 2

Treatment	Variation of customer’s appearance and amount of candy
<i>Rich-Large</i>	High income customer buys large amount (150 grams)
<i>Rich-Small</i>	High income customer buys small amount (50 grams)
<i>Poor-Large</i>	Low income customer buys large amount (150 grams)
<i>Poor-Small</i>	Low income customer buys small amount (50 grams)

Note: The experimental treatments are implemented in each kiosk with hidden scale ($n=50$) through four consecutive purchases.

Experiment 2.

Again, talking was held neutral and constant over all purchases. We visited 50 kiosks that rely on hidden scales.⁸ Our design was exploratory in the sense that good theoretical arguments exist to justify predictions that either highlight the intrapersonal stability of fraudulent behavior or the situational aspect of the income-perception or amount bought (e.g., Balafoutas et al., 2013). By varying the amounts bought and the physical appearance of the buyer, we designed a small panel that assigns four observations to each seller, therefore allowing a test of the behavioral consistency of sellers over time.

5.4.1 Procedure

Each kiosk was visited by four different buyers within one hour on an evening in late November 2011.⁹ Between each sale, 15 minute-breaks ensured that no suspicion was raised by sellers about repeated sales of largely identical candy bags. Similar to Experiment 1, the buyers entered the kiosk and collected either 50 or 150 grams of candy (i.e., roughly 5 or 15 pieces).

5.4.2 Results and Discussion

Experiment 2 addresses the robustness of the results by implementing a panel in the field while also varying perceptions of income and amounts of candy bought. Table 5.3 displays the mean of all Δ_i and summarizes

⁸We were able to identify 5 additional kiosks between Experiment 1 and 2 that were previously missed in scanning the city for kiosks.

⁹Each buyer kept his role during the entire experiment, e.g., the same buyer in the *rich* condition always bought a *large* amount of candy.

Table 5.3: Overview of results from Experiments 1 and 2

Treatment	n	Δ_i (means)	% honest	% overcharging	% undercharging
Exp. 1: <i>Hidden Scale</i>	45	1.69	9%	40%	51%
Exp. 1: <i>Public Scale</i>	37	-4	43%	8%	49%
Exp. 2: <i>Rich-Large</i>	50	1.42	10%	44%	46%
Exp. 2: <i>Rich-Small</i>	50	-0.36	8%	32%	60%
Exp. 2: <i>Poor-Large</i>	50	-0.1	20%	38%	42%
Exp. 2: <i>Poor-Small</i>	50	1.02	16%	38%	46%

Notes: Δ_i denotes the difference between the price paid and the actual price. In Experiment 2 scales were always hidden. *Rich* and *poor* denotes the appearance of the buyer. *Large* and *small* specifies the quantity of candy bought.

the frequencies of honesty, over-, and undercharging.¹⁰ Pairwise comparisons of the mean Δ_i between the treatments of Experiment 2 show no significant differences. As shown in Table 5.3, the largest mean of Δ_i can be observed in treatment *Rich-Large*. Also most overcharging took place in situations where the buyer appears to be rich and purchases a large amount of candy. However, frequencies of honesty, overcharging, and undercharging are not statistically different between the treatments in Experiment 2.¹¹

Table 5.4 shows a series of linear regression models that include dummy variables for the treatments in Experiment 2. In models (1)-(3), the de-

¹⁰The average size of overcharging (conditional on overcharging) is 10 cents in *Rich-Large*, 6 cents in *Rich-Small*, 5 cents in *Poor-Large*, and 6 cents in *Poor-Small* (see also Table A5.2 in the appendix). The average size of undercharging (conditional on undercharging) is 5 cents in *Rich-Large*, 4 cents in *Rich-Small*, 5 cents in *Poor-Large*, and 3 cents in *Poor-Small* (see also Table A5.3 in the appendix).

¹¹Pairwise comparisons of Δ_i in the *Public Scale* treatment with the Δ_i of the treatments in Experiment 2 reveal that the differences between the prices paid and the actual prices in all treatments of Experiment 2 (except treatment *Rich-Small*) are significantly larger (all p -values < 0.03, one-sided Fisher-Pitman permutation test). Pairwise comparisons of the frequencies of honesty and overcharging between *Public Scale* and the treatments of Experiment 2 all reveal significant differences (all p -values < 0.05, one-sided χ^2 -tests), i.e., more overcharging and less honesty is observed in all treatments of Experiment 2 compared to the *Public Scale* treatment in Experiment 1. The frequencies of undercharging are not different between treatment *Public Scale* and the treatments of Experiment 2. See Table A5.6 in the appendix.

Table 5.4: OLS regressions on Δ_i for Experiments 1 and 2

Dependent Variable:	Size of Δ_i			Size of overcharging (conditional on overcharging)	
	(1)	(2)	(3)	(4)	(5)
<i>Rich-Large</i>	0.400 (1.846)		0.400 (1.839)	4.030 (2.796)	
<i>Rich-Small</i>	-1.380* (0.726)		-1.38 (0.723)	0.003 (1.361)	
<i>Poor-Large</i>	-1.120 (0.928)		-1.120 (0.926)	-1.316 (1.485)	
<i>Public Scale</i>			-5.013*** (1.805)		
Status (1 if rich)		-1.380* (0.726)			0.003 (1.361)
Quantity (1 if large)		-1.120 (0.928)			-1.316 (1.485)
Interaction: Status×Quantity		2.900 (2.008)			5.343* (2.914)
Gender (1 if female)	-0.217 (1.956)	-0.217 (1.956)	0.390 (1.743)	-1.714 (2.823)	-1.714 (2.823)
Constant	1.042 (0.772)	1.042 (0.772)	1.02 (0.746)	6.444*** (0.999)	6.444*** (0.999)
Observations	200	200	237	76	76
R-squared	0.008	0.008	0.042	0.083	0.083

Notes: OLS-regression coefficients (robust standard errors clustered on kiosk-level in parentheses), reference group: *Poor-Small*. *Rich-Large*, *Rich-Poor*, *Poor-Large*, and *Public Scale* represent dummy variables for the respective treatments. Significance at the 1, 5, and 10 percent level is denoted by ***, **, and *, respectively.

pendent variable is Δ_i , and it can be shown that the treatment dummies of Experiment 2 have almost no effect on Δ_i (except for treatment *Rich-Small*). In model (3), a dummy variable for treatment *Public Scale* is included, which has a significant negative effect on Δ_i . In models (4)-(5), the dependent variable is the size of overcharging conditional on overcharging. In these models, the treatment dummies from Experiment 2 also have no significant effect.

As can be seen in Table 5.4, we find no strong isolated effects of either the visual status appearance or the quantity of candy bought, whereas the interaction of the two variables is weakly significant (model 5), suggesting that sellers charge buyers more if they have a high status and purchase large quantities. All models include clustered standard errors on the

Table 5.5: Classification of predominant seller types

Type	Proportion
Honest Type	48% (0%)
Liar Type	16% (26%)
Generous Type	4% (38%)
Not classifiable	32% (36%)

Note: Numbers in brackets shows the classification of seller types employing the less conservative definitions of over- and undercharging.

kiosk-level, based on the fact that each seller has made several sales. Thus, when buying large amounts, customers in the high income role are more often overcharged than customers in low income role.

Next, we analyze the robustness of seller behavior in kiosks with hidden scales in Experiment 2. The four repeated purchases enable measurement of behavioral consistency by sellers. To achieve this, we relied on the following conservative protocol. We define three types of sellers, i.e., the predominantly honest, liar, and generous type. A seller is classified as predominantly honest (liar, generous) type if he behaves at least in three of the four sales honestly (fraudulently, generously). Honest (fraudulent, generous) is defined such that in three of the four sales the charged price must fall within the target range, i.e., derivations of less of than 5 cents lead to being categorized as honest, deviations greater than +4 lead to a categorization as fraudulent, deviations smaller than -4 cents lead to a categorization as generous.¹²

As Table 5.5 reveals, 68% of the sellers can be classified as one of the three types. 48% of the sellers behave consistently honest. 16% of the sellers systematically defraud their customers, and 4% of the sellers can be classified as consistently generous. Hence, it can be assumed that the sellers who cannot be classified (32%) somehow react to the treatment variations in Experiment 2.

The numbers in brackets in Table 5.5 show the type classification when employing a less conservative type definition, i.e., a seller is classified as honest (fraudulent, generous) if at least in three of the four sales the charged price is correct (is too large, is too small). In this case, 26%

¹²We define overcharging (undercharging) this way because cent-prices are often rounded up or down to an amount that ends at a five cents point. The method follows protocols similar to "Swedish Rounding", where the price is rounded to the a specific coin, e.g., the five cents coin. Hence, we gave all sellers the benefit of the doubt that they did not utilize coins smaller than 5 cents.

of the sellers can be classified as liar types and 38% as generous types. However, none of the sellers can be classified as an honest type, which suggests that rounding seems to be prevalent to some extent. Therefore, we employed the above mentioned more conservative definition of overcharging (undercharging) for the complete analysis which is defined by demanding at least five cents above (below) the actual price from the customer, i.e., only if the price paid is more than four cents above (below) the correct price a customer is – according to this definition – overcharged (undercharged). By employing this definition, the results remain robust (see numbers in brackets Table A5.6 in the appendix for an application of this definition).

5.5 General Discussion

The present research addressed the occurrence and behavioral consistency of overcharging in a market that shares similar characteristics with a credence goods market, but rules out overtreatment. Additionally, the research explored discrimination against particular buyers in such markets. Drawing on the results from two field experiments, we replicate the core result of credence good research that customers are harmed when informational asymmetries permit sellers to do this. We find that when market characteristics invite overcharging customers face a risk of paying too much, a risk that is virtually eliminated when the markets do not have this credence good attribute. Importantly, overcharging cannot be mistaken for potentially prosocial overtreatment due to best intentions. In our context, overcharging is either erroneous or malicious (although we carefully allow rounding to see how robust overcharging is). In addition, we find that high status customers who purchase large quantities are particularly prone to being overcharged. This highlights some discrimination against people who might be able to afford a small "price-premium", especially when the likelihood of being detected is small. Furthermore, many sellers behave quite consistently across sales, delivering first experimental evidence for the question of behavioral consistency. Interestingly, the majority of sellers can be classified as honest type. Drawing on our research, it seems that certain bad apples spoil the barrel. While lay-wisdom might quip, for instance, that "all car mechanics are frauds", our

data suggest that only a minority of expert sellers cheats, although quite systematically.

Therefore, our results map into the existing literature highlighting the risks for customers that are also found in credence goods markets. Several theoretical and empirical analyzes have highlighted the occurrence of fraudulent behavior in markets for credence goods (see, Darby and Karni, 1973; Dulleck et al., 2011). Building on previous research (e.g., Balafoutas et al., 2013), we show the negative effects of these market characteristics for customers, but we are also able to show in a real-world market setting that only a small proportion of sellers systematically overcharge their customers when they have an easy possibility to defraud, i.e., customers cannot see the display of a scale on which candy is weighed to determine the price. Although the profits from fraud are quite small, more than 15% of sellers systematically betray their customers, yielding a small, but steady profit. While some research emphasizes institutional reasons as a driver of overtreatment bias (e.g., Beck et al., 2014), overcharging could likely be addressed in an easier way. Our context shows that transparency of prices completely mitigates the risk for customers. In the case of openly visible scales, overcharging is negligible.

With our method of buying from the same individuals repeatedly, we can answer some questions regarding behavioral types. However, our quasi-experimental design in Experiment 1 lacks one key feature in field experiments - randomization (see, Harrison and List, 2004). However, we only utilize Experiment 1 to show that overcharging occurs in kiosks, while our main research question addressing behavioral consistency and discrimination against high status persons is solely addressed in Experiment 2. To address the lack of randomization, we can show that both populations with public and hidden scale, respectively, do not systematically differ regarding demographics and other personality traits. Another potential shortcoming is that the financial consequences for people involved are rather low. However, we can deliver an insight into a real market and show that fraudulent as well as generous behavior might occur. This finding might be transferable to other markets with higher financial stakes.

Further experimental research is needed to better understand behavioral patterns across time in other markets for credence goods. It would

be particularly interesting to address the stability of overcharging in professional services firms that sell services that are highly prone to fraud and highly sensitive, e.g., programming or data-services. In fact, trustworthiness of professional services firms is a key issue that becomes more important the more sensitive the market becomes. Therefore, economists need to gather detailed knowledge about the exact mechanisms that certain institutional details impose on market outcomes.

5.6 Appendix

Figure A5.1: Map of kiosks



Note: Bright kiosks use hidden scale, whereas dark kiosks apply public scales.

Table A5.1: Overview of sellers' demographics and personality traits

Question on	<i>Hidden Scale</i> (n=27)	<i>Public Scale</i> (n=10)	Difference
Age	39.2	37.4	1.8
Gender	75% male	83% male	-8%
Foreign	71%	91%	-20%
Openness	2.82	2.5	0.32
Conscientiousness	2.28	2.33	-0.05
Extraversion	3.43	3.08	0.35
Agreeableness	1.83	2.08	-0.25
Neuroticism	1.81	1.63	0.18

Note: Participation was not incentivized. Differences of all personality traits between *Hidden Scale* and *Public Scale* are statistically not different (all p -values > 0.2 , MWU-tests, two-sided). Scale reliabilities of the traits are: Openness (Cronbach's $\alpha=0.374$), Conscientiousness ($\alpha=0.439$), Extraversion ($\alpha=0.204$), Agreeableness ($\alpha=0.453$), Neuroticism ($\alpha=0.372$).

Table A5.2: Amount of overcharging in cents

Treatment	n	Amount of overcharging	% of actual price overcharged	n (overcharging)
Exp. 1: <i>Hidden Scale</i>	45	12 (24)	8% (18%)	18 (8)
Exp. 1: <i>Public Scale</i>	37	5 (9)	4% (6%)	3 (1)
Exp. 2: <i>Rich-Large</i>	50	10 (17)	6% (10%)	22 (12)
Exp. 2: <i>Rich-Small</i>	50	6 (10)	13% (22%)	16 (8)
Exp. 2: <i>Poor-Large</i>	50	5 (11)	3% (7%)	19 (7)
Exp. 2: <i>Poor-Small</i>	50	6 (10)	14% (25%)	19 (9)
Total	282	8 (15)	7% (14%)	97 (45)

Notes: Amounts of overcharging in cents are conditional on overcharging. Numbers in brackets show the amounts employing the more conservative definition of overcharging, i.e., overcharging is defined by demanding at least five cents above the actual price from the customer.

Table A5.3: Amount of undercharging in cents

Treatment	n	Amount of undercharging	% of actual price undercharged	n (undercharging)
Exp. 1: <i>Hidden Scale</i>	45	6 (12)	4% (8%)	23 (9)
Exp. 1: <i>Public Scale</i>	37	9 (17)	5% (9%)	18 (8)
Exp. 2: <i>Rich-Large</i>	50	5 (8)	3% (5%)	22 (9)
Exp. 2: <i>Rich-Small</i>	50	4 (8)	9% (18%)	30 (9)
Exp. 2: <i>Poor-Large</i>	50	5 (8)	3% (5%)	21 (10)
Exp. 2: <i>Poor-Small</i>	50	3 (6)	7% (14%)	23 (5)
Total	282	5 (10)	4% (8%)	137 (50)

Notes: Amounts of undercharging in cents are conditional on undercharging. Numbers in brackets show the amounts employing the more conservative definition of undercharging, i.e., undercharging is defined by demanding at least five cents below the actual price from the customer.

Table A5.4: OLS regressions for Experiment 1

Dependent Variable:	Size of Δ_i				
	(1)	(2)	(3)	(4)	(5)
<i>Hidden Scale</i>	5.689** (2.675)	5.491** (2.552)	3.436 (3.954)	1.327 (3.945)	
Gender (1 if female)		6.604 (7.891)	5.806 (8.784)	5.03 (8.596)	
Age			0.265 (0.138)	0.211 (0.159)	
Openness				2.029 (1.716)	
Conscientiousness				-1.078 (2.277)	
Extraversion				2.615 (1.897)	
Agreeableness				0.548 (2.004)	
Neuroticism				-1.194 (2.876)	
Fill in Survey (1 if yes)					5.343* (2.764)
Constant	-4** (1.644)	-4.535** (1.814)	-11.811* (6.691)	-18.711 (11.686)	-3.289* (2.019)
Observations	82	82	37	37	82
<i>R</i> -squared	0.05	0.074	0.131	0.131	0.044

Notes: OLS-regression coefficients (robust standard errors in parentheses), reference group: *Public Scale*. *Hidden Scale* represents a dummy variable for this treatment. Numbers of observations drop in model (3) and (4) because of missing values. In model (5), “Fill in Survey” represents a dummy variable for sellers’ willingness to fill in the questionnaire. Significance at the 1, 5, and 10 percent level is denoted by ***, **, and *, respectively.

Table A5.5: Weight estimation of candy bags in laboratory

Weight of candy bag	<i>n</i>	Estimation (mean in grams)	Standard deviation
50 grams	32	75.38	29.39
150 grams	32	218.41	64.78

Note: Accuracy of estimation was incentivized.

Table A5.6: Results from Experiments 1 and 2

Treatment	n	Mean Δ_i	% honest	% overcharging	% undercharging
Exp. 1: <i>Hidden Scale</i>	45	1.69 (1.84) $\sqrt{***}$	9% (60%) $\wedge+++(+++)$	40% (18%) $\sqrt{+++}(++)$	51% (20%) \vee
Exp. 1: <i>Public Scale</i>	37	-4 (-3.41)	43% (78%)	8% (3%)	49% (22%)
Exp. 2: <i>Rich-Large</i>	50	1.42 (1.6) $\sqrt{***}$	10% (58%) $\wedge+++(+++)$	44% (24%) $\sqrt{+++}(+++)$	46% (20%) \wedge
Exp. 1: <i>Public Scale</i>	37	-4 (-3.41) $\wedge***$	43% (78%) $\sqrt{+++}(+++)$	8% (3%) $\wedge+++(+++)$	49% (22%) \wedge
Exp. 2: <i>Rich-Small</i>	50	-0.36 (0.24)	8% (66%)	32% (16%)	60% (18%)
Exp. 2: <i>Poor-Large</i>	50	-0.1 (-0.06) $\sqrt{***}$	20% (66%) $\wedge++(+++)$	38% (14%) $\sqrt{+++}(++)$	42% (20%) \wedge
Exp. 1: <i>Public Scale</i>	37	-4 (-3.41) $\wedge****$	43% (78%) $\sqrt{+++}(+++)$	8% (3%) $\wedge+++(+++)$	49% (22%) \vee
Exp. 2: <i>Poor-Small</i>	50	1.02 (1.24)	16% (72%)	38% (18%)	46% (10%)

Notes: Δ denotes the difference between the price paid and the actual price. In Experiment 2 scales were always hidden. *Rich* and *poor* denotes the appearance of the buyer. *Large* and *small* specifies the quantity of candy bought. n stands for the number of observations. Stars display the significance of a one-sided Fisher-Pitman permutation test (*=10 %-level, **=5 %-level, ***=1 %-level). Plus and minus signs display the significance of a one-sided χ^2 -test ($-$ (+)=10 %-level, $--$ (++)=5 %-level, $---$ (+++)=1 %-level). Numbers in brackets show the results and levels of significance employing the more conservative definition of overcharging and undercharging, i.e., overcharging (undercharging) is defined by demanding at least five cents above (below) the actual price from the customer.

Bibliography

- ABELER, J., A. BECKER, AND A. FALK (2014): “Representative evidence on lying costs,” *Journal of Public Economics*, 113, 96 – 104.
- ANDREONI, J. AND B. D. BERNHEIM (2009): “Social image and the 50-50 norm: A theoretical and experimental analysis of audience effects,” *Econometrica*, 77, 1607–1636.
- BALAFOUTAS, L., A. BECK, R. KERSCHBAMER, AND M. SUTTER (2013): “What drives taxi drivers? A field experiment on fraud in a market for credence goods,” *The Review of Economic Studies*, 80, 876–891.
- BANDURA, A. (1991): “Social cognitive theory of moral thought and action,” in *Handbook of Moral Behavior and Development*, ed. by W. Kurtines and J. Gerwitz, Hillsdale, NJ: Erlbaum, vol. 1, 45–103.
- BANDURA, A., C. BARBARANELLI, G. V. CAPRARA, AND C. PASTORELLI (1996): “Mechanisms of moral disengagement in the exercise of moral agency,” *Journal of Personality and Social Psychology*, 71, 364–374.
- BAZERMAN, M. H. AND F. GINO (2012): “Behavioral ethics: Toward a deeper understanding of moral judgment and dishonesty,” *Annual Review of Law and Social Science*, 8, 85–104.
- BAZERMAN, M. H. AND A. E. TENBRUNSEL (2011): *Blind spots: Why we fail to do what’s right and what to do about it*, Princeton University Press.
- BECK, A., R. KERSCHBAMER, J. QIU, AND M. SUTTER (2014): “Car mechanics in the lab: Investigating the behavior of real experts on experimental markets for credence goods,” *European University Institute Economics Working Papers ECO2014/02*.

- BECKER, G. S. (1968): “Crime and punishment: An economic approach,” *Journal of Political Economy*, 76, 169.
- BEISECKER, T., G. WALKER, AND J. BART (1989): “Knowledge versus ignorance as bargaining strategies: The impact of knowledge about other’s information level,” *Social Science Journal*, 26, 161–172.
- BELOT, M. AND M. SCHRÖDER (2012): “Sloppy work, lies and theft: A novel experimental design to study counterproductive behavior,” *Otto von Guericke University Magdeburg Working Paper*, 18.
- BÉNABOU, R. AND J. TIROLE (2002): “Self-confidence and personal motivation,” *The Quarterly Journal of Economics*, 117, 871–915.
- (2006): “Incentives and prosocial behavior,” *The American Economic Review*, 96, 1652–1678.
- BERGER, J., C. HERBERTZ, AND D. SLIWKA (2010): “Incentives and cooperation in firms: Field evidence,” *IZA Discussion Papers*, 5618.
- BICCHIERI, C. (2006): *The grammar of society: The nature and dynamics of social norms*, Cambridge University Press.
- BICCHIERI, C. AND A. LEV-ON (2007): “Computer-mediated communication and cooperation in social dilemmas: An experimental analysis,” *Politics, Philosophy & Economics*, 6, 139–168.
- (2011): “Studying the ethical implications of e-trust in the lab,” *Ethics and information technology*, 13, 5–15.
- BICCHIERI, C. AND E. XIAO (2009): “Do the right thing: But only if others do so,” *Journal of Behavioral Decision Making*, 22, 191–208.
- BIEL, A. AND J. THOGERSEN (2007): “Activation of social norms in social dilemmas: A review of the evidence and reflections on the implications for environmental behaviour,” *Journal of Economic Psychology*, 28, 93–112.
- BLOUNT, S. (1995): “When social outcomes aren’t fair: The effect of causal attributions on preferences,” *Organizational Behavior and Human Decision Processes*, 63, 131–144.

- BOHNET, I. AND B. S. FREY (1999): "Social distance and other-regarding behavior in dictator games: Comment," *The American Economic Review*, 89, 335–339.
- BOK, S. (1999): *Lying - Moral choice in public and private life*, Vintage Books.
- BOLTON, G. E. AND A. OCKENFELS (2000): "ERC: A theory of equity, reciprocity, and competition," *The American Economic Review*, 90, 166–193.
- BOLTON, G. E. AND R. ZWICK (1995): "Anonymity versus punishment in ultimatum bargaining," *Games and Economic Behavior*, 10, 95–121.
- BORDIA, P. (1997): "Face-to-face versus computer-mediated communication: A synthesis of the experimental literature," *Journal of Business Communication*, 34, 99–118.
- BRANDTS, J. AND G. CHARNES (2011): "The strategy versus the direct-response method: A first survey of experimental comparisons," *Experimental Economics*, 14, 375–398.
- BROBERG, T., T. ELLINGSEN, AND M. JOHANNESSON (2007): "Is generosity involuntary?" *Economics Letters*, 94, 32–37.
- BROSIG, J. (2002): "Identifying cooperative behavior: Some experimental results in a prisoner's dilemma game," *Journal of Economic Behavior & Organization*, 47, 275–290.
- BROSIG, J., J. WEIMANN, AND A. OCKENFELS (2003): "The effect of communication media on cooperation," *German Economic Review*, 4, 217–241.
- BUCCIOL, A. AND M. PIOVESAN (2011): "Luck or cheating? A field experiment on honesty with children," *Journal of Economic Psychology*, 32, 73–78.
- BUSS, A. H. AND S. R. BRIGGS (1984): "Drama and the self in social interaction," *Journal of Personality and Social Psychology*, 47, 1310–1324.

- CADSBY, C. B., F. SONG, AND F. TAPON (2010): “Are you paying your employees to cheat? An experimental investigation,” *BE Journal of Economic Analysis & Policy*, 10, 1–30.
- CAMERER, C. F. (2003): *Behavioral game theory*, Princeton University Press.
- CHARNESS, G. AND U. GNEEZY (2008): “What’s in a name? Anonymity and social distance in dictator and ultimatum games,” *Journal of Economic Behavior & Organization*, 68, 29–35.
- CHARNESS, G., E. HARUVY, AND D. SONSINO (2007): “Social distance and reciprocity: An internet experiment,” *Journal of Economic Behavior & Organization*, 63, 88–103.
- CHARNESS, G. AND M. RABIN (2002): “Understanding social preferences with simple tests,” *The Quarterly Journal of Economics*, 117, 817–869.
- CONRADS, J. AND B. IRLLENBUSCH (2013): “Strategic ignorance in ultimatum bargaining,” *Journal of Economic Behavior & Organization*, 92, 104–115.
- CONRADS, J., B. IRLLENBUSCH, R. M. RILKE, A. SCHIELKE, AND G. WALKOWITZ (2014): “Honesty in tournaments,” *Economics Letters*, 123, 90 – 93.
- CONRADS, J., B. IRLLENBUSCH, R. M. RILKE, AND G. WALKOWITZ (2013): “Lying and team incentives,” *Journal of Economic Psychology*, 34, 1 – 7.
- CRAWFORD, V. P. AND J. SOBEL (1982): “Strategic information transmission,” *Econometrica*, 1431–1451.
- CROSON, R. T. A. (2005): “Deception in economics experiments,” in *Deception in Markets An Economic Analysis*, ed. by C. Gerschlager, Palgrave Macmillan, chap. 5, 113–130.
- DAFT, R. L. AND R. H. LENGEL (1986): “Organizational information requirements, media richness and structural design,” *Management Science*, 32, 554–571.

- DANA, J., D. M. CAIN, AND R. M. DAWES (2006): “What you don’t know won’t hurt me: Costly (but quiet) exit in dictator games,” *Organizational Behavior and Human Decision Processes*, 100, 193–201.
- DANA, J., R. A. WEBER, AND J. X. KUANG (2007): “Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness,” *Economic Theory*, 33, 67–80.
- DANILOV, A., T. BIEMANN, T. KRING, AND D. SLIWKA (2013): “The dark side of team incentives: Experimental evidence on advice quality from financial service professionals,” *Journal of Economic Behavior & Organization*, 93, 266–272.
- DARBY, M. R. AND E. KARNI (1973): “Free competition and the optimal amount of fraud,” *Journal of Law and Economics*, 16, 67–88.
- DEPAULO, B. M., D. A. KASHY, S. E. KIRKENDOL, M. M. WYER, AND J. A. EPSTEIN (1996): “Lying in everyday life,” *Journal of Personality and Social Psychology*, 70, 979–995.
- DOMENIGHETTI, G., A. CASABIANCA, F. GUTZWILLER, AND S. MARTINOLI (1993): “Revisiting the most informed consumer of surgical services: The physician-patient,” *International Journal of Technology Assessment in Health Care*, 9, 505–513.
- DREBER, A. AND M. JOHANNESSON (2008): “Gender differences in deception,” *Economics Letters*, 99, 197–199.
- DULLECK, U. AND R. KERSCHBAMER (2006): “On doctors, mechanics, and computer specialists: The economics of credence goods,” *Journal of Economic Literature*, 5–42.
- DULLECK, U., R. KERSCHBAMER, AND M. SUTTER (2011): “The economics of credence goods: An experiment on the role of liability, verifiability, reputation, and competition,” *The American Economic Review*, 101, 526–555.
- EBELING, F., C. FELDHAUS, AND J. FENDRICH (2012): “Follow the leader or follow anyone—Evidence from a natural field experiment,” *Cologne Graduate School Working Paper*.

- ELLINGSEN, T., M. JOHANNESSON, J. LILJA, AND H. ZETTERQVIST (2009): “Trust and truth,” *The Economic Journal*, 119, 252–276.
- FALK, A. AND E. FEHR (2003): “Why labour market experiments?” *Labour Economics*, 10, 399–406.
- FALK, A., E. FEHR, AND U. FISCHBACHER (2003): “On the nature of fair behavior,” *Economic Inquiry*, 41, 20–26.
- (2008): “Testing theories of fairness—Intentions matter,” *Games and Economic Behavior*, 62, 287–303.
- FALK, A. AND J. J. HECKMAN (2009): “Lab experiments are a major source of knowledge in the social sciences,” *Science*, 326, 535–538.
- FALK, A. AND N. SZECH (2013): “Morals and markets,” *Science*, 340, 707–711.
- FEHR, E. AND K. M. SCHMIDT (1999): “A theory of fairness, competition, and cooperation,” *The Quarterly Journal of Economics*, 114, 817.
- FELDMAN, R. (2009): *Liar: The truth about lying*, Virgin Books, Grand Central Publishing, New York.
- FISCHBACHER, U. (2007): “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 10, 171–178.
- 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.
- FISCHBACHER, U. AND V. UTIKAL (2013): “On the acceptance of apologies,” *Games and Economic Behavior*, 82, 592–608.
- FISHER, R. AND W. L. URY (1991): *Getting to Yes: Negotiating agreement without giving in*, Penguin, 2nd edition ed.
- FRANK, R. H. (1988): *Passions within reason: The strategic role of the emotions.*, WW Norton & Co.
- GÄCHTER, S. AND E. FEHR (1999): “Collective action as a social exchange,” *Journal of Economic Behavior & Organization*, 39, 341–369.

- GEHRIG, T., W. GÜTH, AND R. LEVÍNSKÝ (2003): “Ultimatum offers and the role of transparency: An experimental study of information acquisition,” *Papers on Strategic Interaction - Max Planck Institute of Economics, Strategic Interaction Group*.
- GEIS, F. L. AND T. H. MOON (1981): “Machiavellianism and deception.” *Journal of Personality and Social Psychology*, 41, 766.
- GIBBONS, R. (1998): “Incentives in organizations,” *Journal of Economic Perspectives*, 12, 115–132.
- GIBSON, R., C. TANNER, AND A. F. WAGNER (2013): “Preference for truthfulness: Heterogeneity among and within individuals,” *The American Economic Review*, 103, 532–548.
- GNEEZY, U. (2005): “Deception: The role of consequences,” *The American Economic Review*, 95, 384–394.
- GNEEZY, U., A. IMAS, AND K. MADARÁSZ (forthcoming): “Conscience accounting: Emotional dynamics and social behavior,” *Management Science*.
- GREENE, J. D. (2012): *The moral brain and how to use it*, Penguin Group, New York.
- GREINER, B. (2003): “An online recruitment system for economic experiments,” in *Forschung und wissenschaftliches Rechnen*, ed. by K. Kremer and V. Macho, GWDG Bericht 63, 79–93.
- GROSSMAN, Z. (2010): “Self-signaling versus social-signaling in giving,” *UCSB Working paper*.
- GROSSMAN, Z. AND J. VAN DER WEELE (2013): “Self-image and strategic ignorance in moral dilemmas,” *UCSB Working paper*.
- GÜTH, W., R. SCHMITTBERGER, AND B. SCHWARZE (1982): “An experimental analysis of ultimatum bargaining,” *Journal of Economic Behavior & Organization*, 3, 367–388.
- GÜTH, W., S. HUCK, AND W. MÜLLER (2001): “The relevance of equal splits in ultimatum games,” *Games and Economic Behavior*, 37, 161–169.

- HAMNER, W. C. AND D. L. HARNETT (1975): “The effects of information and aspiration level on bargaining behavior,” *Journal of Experimental Social Psychology*, 11, 329–342.
- HANCOCK, J. T., J. THOM-SANTELLI, AND T. RITCHIE (2004): “Deception and design: The impact of communication technology on lying behavior,” in *Proceedings of the SIGCHI conference on human factors in computing systems*, ACM, 129–134.
- HARBRING, C. AND B. IRLENBUSCH (2011): “Sabotage in tournaments: Evidence from a laboratory experiment,” *Management Science*, 57, 611–627.
- HARRISON, G. W. AND J. A. LIST (2004): “Field experiments,” *Journal of Economic Literature*, 1009–1055.
- HOFFMAN, E., K. MCCABE, AND V. L. SMITH (1996): “Social distance and other-regarding behavior in dictator games,” *The American Economic Review*, 86, 653–660.
- HOUSER, D., S. VETTER, AND J. WINTER (2012): “Fairness and cheating,” *European Economic Review*, 56, 1645–1655.
- KAGEL, J. H., C. KIM, AND D. MOSER (1996): “Fairness in ultimatum games with asymmetric information and asymmetric payoffs,” *Games and Economic Behavior*, 13, 100–110.
- KARTIK, N. (2009): “Strategic communication with lying costs,” *Review of Economic Studies*, 76, 1359–1395.
- KASHY, D. A. AND B. M. DEPAULO (1996): “Who lies?” *Journal of Personality and Social Psychology*, 70, 1037–1051.
- KERSCHBAMER, R., M. SUTTER, AND U. DULLECK (2009): “The impact of distributional preferences on (experimental) markets for expert services,” *IZA Discussion Paper No. 4647*.
- KNOBE, J. AND S. NICHOLS (2008): *Experimental philosophy*, Oxford University Press.
- LAZEAR, E. P. AND M. GIBBS (2009): *Personnel economics in practice*, Wiley.

- LEDFOORD, G. E., E. E. LAWLER, AND S. A. MOHRMAN (1995): "Reward innovations in fortune 1000 companies," *Compensation & Benefits Review*, 27, 76–80.
- LEVINE, D. K. (1998): "Modeling altruism and spitefulness in experiments," *Review of economic dynamics*, 1, 593–622.
- LEWIS, M. (1993): "The development of deception," in *Lying and deception in everyday life*, ed. by M. Lewis and C. Saarni, New York: The Guilford Press, 90–105.
- LIST, J. (2006): "Field experiments: A bridge between lab and naturally occurring data," *The BE Journal of Economic Analysis & Policy*, 6.
- LÓPEZ-PÉREZ, R. (2012): "The power of words: A model of honesty and fairness," *Journal of Economic Psychology*, 33, 642–658.
- MAZAR, N. AND P. AGGARWAL (2011): "Greasing the palm: Can collectivism promote bribery?" *Psychological Science*, 22, 843–849.
- MAZAR, N., O. AMIR, AND D. ARIELY (2008): "The dishonesty of honest people: A theory of self-concept maintenance," *Journal of Marketing Research*, 45, 633–644.
- MAZAR, N. AND D. ARIELY (2006): "Dishonesty in everyday life and its policy implications," *Journal of Public Policy & Marketing*, 25, 117–126.
- MILGRAM, S. (1965): "Some conditions of obedience and disobedience to authority," *Human relations*, 18, 57–76.
- MONIN, B. AND D. T. MILLER (2001): "Moral credentials and the expression of prejudice." *Journal of Personality and Social Psychology*, 81, 33.
- NALBANTIAN, H. R. AND A. SCHOTTER (1997): "Productivity under group incentives: An experimental study," *The American Economic Review*, 87, 314–341.
- POULSEN, A. U. AND M. W. M. ROOS (2010): "Do people make strategic commitments? Experimental evidence on strategic information avoidance," *Experimental Economics*, 13, 206–225.

- POULSEN, A. U. AND J. H. W. TAN (2007): “Information acquisition in the ultimatum game: An experimental study,” *Experimental Economics*, 10, 391–409.
- PROCTOR, R. AND L. L. SCHIEBINGER (2008): *Agnotology: The making and unmaking of ignorance*, Stanford University Press.
- PRUCKNER, G. J. AND R. SAUSGRUBER (2013): “Honesty on the streets: A field study on newspaper purchasing,” *Journal of the European Economic Association*, 11, 661–679.
- RAMMSTEDT, B. AND O. P. JOHN (2007): “Measuring personality in one minute or less: A 10-item short version of the big five inventory in English and German,” *Journal of Research in Personality*, 41, 203–212.
- ROSS, W. T. AND D. C. ROBERTSON (2000): “Lying: The impact of decision context,” *Business Ethics Quarterly*, 10, 409–440.
- ROTH, A. E. (1995): “Bargaining experiments,” in *Handbook of Experimental Economics*, ed. by J. Kagel and A. E. Roth, Princeton Univ. Press, Princeton, 253–348.
- ROTH, A. E. AND J. K. MURNIGHAN (1982): “The role of information in bargaining: An experimental study,” *Econometrica*, 50, 1123–1142.
- ROWATT, W. C., M. R. CUNNINGHAM, AND P. B. DRUEN (1998): “Deception to get a date,” *Personality and Social Psychology Bulletin*, 24, 1228–1242.
- SANDEL, M. J. (2012): *What money can't buy: The moral limits of markets*, Macmillan.
- (2013): “Market reasoning as moral reasoning: Why economists should re-engage with political philosophy,” *The Journal of Economic Perspectives*, 27, 121–140.
- SHELLING, T. C. (1956): “An essay on bargaining,” *The American Economic Review*, 46, 281–306.
- (1960): *The strategy of conflict*, Harvard University Press.

- SCHNEIDER, H. S. (2012): “Agency problems and reputation in expert services: Evidence from auto repair,” *The Journal of Industrial Economics*, 60, 406–433.
- SCHRAM, A. AND G. CHARNES (2011): “Social and moral norms in the laboratory,” *UCSB Working Paper*.
- SCHWARTZ, S. H. (1992): “Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries,” *Advances in experimental social psychology*, 25, 1–65.
- SCHWEITZER, M. E. AND C. K. HSEE (2002): “Stretching the truth: Elastic justification and motivated communication of uncertain information,” *Journal of Risk and Uncertainty*, 25, 185–201.
- SCHWEITZER, M. E., L. ORDONEZ, AND B. DOUMA (2004): “Goal setting as a motivator of unethical behavior,” *Academy of Management Journal*, 47, 422–432.
- SEITHE, M. (2012): “Introducing the Bonn experiment system,” *Bonn Econ Discussion Papers, No. 01/2012*.
- SELTEN, R. (1967): “Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopol-experiments,” in *Beiträge zur experimentellen Wirtschaftsforschung*, ed. by H. Sauermann, Tuebingen: J.C.B. Mohr, 136–168.
- SIEGEL, S. AND L. E. FOURAKER (1960): *Bargaining and group decision making: Experiments in bilateral monopoly.*, New York: McGraw-Hill.
- SMITH, A. (2010): *The theory of moral sentiments*, Penguin.
- SUTTER, M. (2009): “Deception through telling the truth?! Experimental evidence from individuals and teams,” *Economic Journal*, 119, 47–60.
- TENBRUNSEL, A. E. AND K. SMITH-CROWE (2008): “Ethical decision making: Where we have been and where we are going,” *The Academy of Management Annals*, 2, 545–607.
- TYLER, J. M. AND R. S. FELDMAN (2004): “Truth, lies, and self-presentation: How gender and anticipated future interaction relate to

- deceptive behavior,” *Journal of Applied Social Psychology*, 34, 2602–2615.
- TYLER, J. M., R. S. FELDMAN, AND A. REICHERT (2006): “The price of deceptive behavior: Disliking and lying to people who lie to us,” *Journal of Experimental Social Psychology*, 42, 69–77.
- VALLEY, K. L., J. MOAG, AND M. H. BAZERMAN (1998): “A matter of trust: Effects of communication on the efficiency and distribution of outcomes,” *Journal of Economic Behavior & Organization*, 34, 211–238.
- VALLEY, K. L., L. THOMPSON, R. GIBBONS, AND M. H. BAZERMAN (2002): “How communication improves efficiency in bargaining games,” *Games and Economic Behavior*, 38, 127–155.
- VRIJ, A. (2010): *Detecting lies and deceit - Pitfalls and opportunities*, John Wiley & Sons.
- WILTERMUTH, S. S. (2011): “Cheating more when the spoils are split,” *Organizational Behavior and Human Decision Processes*, 115, 157–168.
- WOLINSKY, A. (1993): “Competition in a market for informed experts’ services,” *RAND Journal of Economics*, 24, 380–380.