On Cooperation, Communication, and Compliance: Essays in Behavioral and Experimental Economics

Marvin Deversi

München 2020

Marvin Deversi: *On Cooperation, Communication, and Compliance: Essays in Behavioral and Experimental Economics*, Dissertation, March 16, 2020

On Cooperation, Communication, and Compliance: Essays in Behavioral and Experimental Economics

Inaugural-Dissertation
zur Erlangung des Grades
Doctor oeconomiae publicae (Dr. oec. publ.)
an der Volkswirtschaftlichen Fakultät
der Ludwig-Maximilians-Universität München

vorgelegt von
Marvin Deversi

2020

Referent:	Prof. Dr. Martin G. Kocher
Koreferent:	Prof. Dr. Florian Englmaier
Promotionsabschlussberatung:	22.07.2020
Berichterstatter:	Prof. Dr. Martin G. Kocher
	Prof. Dr. Florian Englmaier
	Prof. Dr. Dirk Sliwka
Mündliche Prüfung:	10.07.2020

Acknowledgements

This dissertation would not have been possible without the help and encouragement of advisers, colleagues, friends, and my family. I am grateful to everyone who in various ways contributed to my research.

First, I would like to thank my adviser Martin Kocher. My dissertation crucially depended on his outstanding support and thoughtful comments. Martin was always extremely reliable and I could approach him with any worries that I had. His guidance and open, down-to-earth nature made being his student a fantastic experience. I am also grateful to Florian Englmaier, my co-adviser. Florian was a great source of encouragement in the preparation process of this dissertation. His confidence in my abilities honored me and at the same time motivated me to work even harder. Our regular research and strategy meetings were both very insightful and extremely enjoyable. I could have not wished for a better mentoring team than Martin and Florian.

I want to thank Dirk Sliwka for agreeing to complete my committee and for providing valuable feedback on my research. I also thank Klaus Schmidt for many inspiring and encouraging conversations. Special thanks go to my co-authors Charles Bellemare, Alessandro Ispano, and Christiane Schwieren. It was great fun working with you. I am also grateful to John List who invited me to the University of Chicago and provided very useful feedback on my work. I owe it to him that I was able to witness Richard Thaler's acceptance speech for receiving the Nobel Prize; an unforgettable experience.

I thank Manuela Firmin and Julia Zimmermann as well as our research assistants for their help, administrative support and enjoyable conversations. I would like to acknowledge the Fulbright Commission for financing my visit to Chicago. Likewise, I thank the Elite Network of Bavaria and the German Science Foundation for their financial support via the doctoral program "Evidence-Based Economics" and the grants CRC TR190 and KO 4100/1-1, respectively. My gratitude also goes to a large number of employees of the company that I worked with for the research presented in Chapters 1 and 2. Without them, conducting this research would not have been possible.

I am especially grateful to Simeon Schudy. Simeon always had an open door and commented on various draft versions of my papers. His efforts did not only profoundly improve this dissertation but, even more importantly, enriched my life at our chair and in Munich tremendously. Simeon became a true mentor, supporter, and friend. I would also like to express my special thanks to Peter Schwardmann. Together with Peter, I started my first project of this dissertation. I am thankful for the trust he put in me and for providing a jump start into the world of research. His input was invaluable for me.

When I arrived in Munich five years ago, I could not nearly imagine how *home* I would feel today. Many friends and colleagues at the Chair for Behavioral and Experimental Economics, the Munich Graduate School of Economics, and the Faculty of Economics contributed to this development. This includes, most of all, my friends Benjamin Häusinger, Christoph Winter, Annemarie Gronau, Tobias Rossmann, Sebastian Lang, Felix Peterhammer, Felix Ehrenfried, Navid Sabet, and Lisa Spantig. A very special thanks goes to Svenja Möller who has always been the most supportive and loyal friend to me.

I am indebted to my parents Cornelia and Rüdiger Deversi and my twin sister Sandra Deversi for their unconditional love that builds the foundation of this dissertation. They made me the person I am today and accepted every of my decisions in life – even if it entailed moving far away from them. Thank you for always being there and giving me the opportunities that were never given to you. Finally, my deepest gratitude goes to Denis Miezal who accompanied me through all the ups and downs in the process of writing this dissertation. He enriches my life in a way that none of my research activities ever will.

Marvin Deversi
Munich in March 2020

Contents

List of Figures

List of Tables

Preface

The neoclassical model of economic behavior is based on the assumption of rationality. It models human decision-making as if individuals maximize their egoistic preferences and correctly form beliefs. In contrast to this approach, the behavioral economics model views human decision-making as neither fully rational nor solely driven by selfish considerations. Instead, it incorporates insights from human psychology such that fallible judgment, social concerns, and mistakes in belief formation play an eminent role (Rabin, 2002).

Behavioral economics originated in the seminal papers of Simon (1955), Kahneman and Tversky (1979), and Thaler (1980). In its infancy, it focused on detecting and describing systematic deviations of human behavior from the neoclassical assumptions. Early work revealed and documented such "behavioral anomalies" in uncontrolled happenstance data. The increasing use of laboratory experiments allowed researchers to identify the causes of such anomalies in controlled environments (Falk and Heckman, 2009). This evidence inspired various models of nonclassical preferences (e.g., Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Laibson, 1997) and beliefs (e.g., Eyster and Rabin, 2005; Kőszegi and Rabin, 2006) that found their way into many subfields of economics previously shaped by the neoclassical paradigm. Among others, behavioral modeling has been applied to game theory (Camerer, 1997), organizational economics (Camerer and Malmendier, 2007), industrial organization (Heidhues and Kőszegi, 2018), or public economics (Bernheim and Taubinsky, 2018).

Nowadays, many of these models are explicitly tested in field environments. Public as well as private organizations all over the world team up with researchers to use insights from behavioral economics models and tackle important real-world problems. Such collaborations range from field experiments on worker incentives (Hossain and List, 2012), performance evaluations (Swift et al., 2013), and price disclosure (Dertwinkel-Kalt et al., 2019) to retirement savings (Thaler and Benartzi, 2004) and tax collection (Hallsworth et al., 2017). They add external validity to phenomena discovered previously in controlled laboratory environments (Levitt and List, 2007; DellaVigna, 2009).

Combining insights from behavioral models and the advantages of experimental methods yields important synergies for researchers as well as policymakers and practitioners. Experimental methods allow decision makers to rigorously evaluate existing policies. Behavioral models provide the structure to put results of such evaluations in perspective and thereby design (and test) appropriate new policies. This process results in evidence-based policy advice (as coined by Roth (1986) researchers can "whisper in the ears of princes" (p. 246)) and allows for a profound discussion of welfare implications of different policies (Chetty, 2015). Further, experimental investigations of policy-oriented research questions yield important indications for necessary refinements of behavioral economics models. For example, evidence from field experiments may help to identify important contextual factors that render behavioral phenomena more or less likely to evolve.

The main reason for the intensifying exchange of behavioral economists with policymakers and practitioners is that they share a common interest in "actual" behavior of humans, may it be students, employees, or other market participants. In this spirit, my dissertation presents four essays that examine the implications of human psychology for economic policy and management practice. All four essays follow a common approach. I combine lessons from behavioral economics models with experimental methods to reveal important insights for the research community as well as for practitioners. Building on the behavioral economics approach, I acknowledge that preferences are not necessarily selfish, that beliefs are often not correctly formed, and that decision-making can be erroneous. Further, I exploit the advantages of, both, laboratory and field experiments.

I use this approach in the context of companies and markets as well as political contexts such as regulation and taxation. In the first two chapters, I look at the extent to which behavioral insights on non-selfish preferences can be applied in companies. Doing so, I investigate two main questions: What are the financial and non-financial consequences of behaving cooperatively in companies? And how can companies promote cooperation among employees? While in the first two chapters I study firms from the inside perspective, in the third chapter I examine communication strategies of firms in markets. I ask how communication needs to be regulated in order to prevent exploitation of consumers that hold wrong beliefs about product quality. In the fourth chapter, I study taxation as another relevant policy domain. I examine the consequences of complex tax filing policies and the appropriation of taxes on compliance of non-rational, non-selfish taxpayers.

In the following, I provide a brief summary of each chapter. The chapters are self-contained and can be read independently. The respective appendices as well as the bibliography can be found at the end of this dissertation.

Chapter 1 is based on joint work with *Martin Kocher* and *Christiane Schwieren*. It starts with the observation that cooperation is vital in most production processes within companies. However, many companies employ incentives or reward institutions that foster selfish behavior (Gratton, 2009). In this chapter, we analyze cooperation within a company setting in order to study the relationship between cooperative attitudes and financial as well as non-financial rewards.

In total, 910 employees of a large software company participate in an incentivized online experiment in which we elicit cooperative attitudes using a modified linear public goods game (based on Fischbacher et al., 2001). We link the experimental measure to record data on financial and non-financial rewards of employees as well as to survey data on other non-financial outcome variables like team cohesion or work satisfaction.

First, we observe high levels of cooperation and the typical conditional contribution patterns in the public goods game. Second, when linking experiment and company data, we find that cooperative attitudes of employees do not pay off in terms of financial rewards within the company. In stark contrast, cooperative employees receive significantly lower wage increases and financial award payments than their more selfish colleagues. Third, as opposed to financial award patterns, cooperative employees receive non-financial benefits such as recognition or friendship as the main reward medium. They also report a higher average work satisfaction.

In contrast to most studies in the experimental laboratory (e.g., Fischbacher and Gächter, 2010), sustained levels of cooperation in our company setting relate to non-financial values of cooperation rather than solely to financial incentives. This suggests that the company achieves high levels of cooperation despite financial disincentives by providing employees with non-financial compensating differentials. Our findings give rise to various management recommendations for companies that want to foster cooperation among their employees.

The evidence in Chapter 1 highlights the importance of non-financial management practices for the cooperative culture in companies. It remains open whether companies can promote cooperation by using dedicated financial incentives. **Chapter 2** addresses this question. Behavioral economists and management scholars have argued that the scope of incentives to increase cooperation in organizations may be limited (e.g., Sliwka, 2007; Bowles and Polanía-Reyes, 2012). One main idea builds on the intuition that conditionally cooperative employees may interpret the provision of incentives as a signal about the prevalence of selfish employees in the organization and thus behave less cooperatively themselves (e.g., Van der Weele, 2012).

I test this hypothesis experimentally using a sample of managers and employees (N = 448) from the software company. In the experiment, employees face a social dilemma situation in which they have a dominant strategy to free-ride on the cooperative efforts of their colleagues. The managers benefit from high cooperation levels among employees and can counter free-riding by setting an incentive that is however costly for them. I exogenously vary whether managers are informed about prevailing cooperation levels among employees, before they can set the incentive to promote cooperation.

Comparing informed versus uninformed incentive choices, the data reveals strong positive effects of incentives that are unaffected by the hypothesized signaling effect. The absence of such effect can be explained by the employees' perception of their managers' intentions. Employees presume that managers choose incentives based on pro-social considerations rather than based on individual profit maximization. Hence, the inference logic that incentive provision indicates low cooperation levels is offset. This mitigating factor has not been explored in the literature so far but entails relevant implications for the optimal design of incentives in organizations. In particular, it gives rise to potential management strategies to prevent the signaling of incentives.

The first two chapters analyze cooperation in the context of firms. **Chapter 3**, which is based on joint work with *Alessandro Ispano* and *Peter Schwardmann*, studies the interaction of firms and consumers in markets. In particular, we investigate why in many market settings unfavorable news are delivered under the disguise of vagueness. An illustrative example of such messages can be found on university websites where universities often refer to the latest US News Ranking as "ranked top 10" rather than as "ranked number 10". These vague messages are not outright lies, but merely put a positive spin on unfavorable news. An open question is whether people are sufficiently naive to be fooled by such positive spin.

We use a theoretical model and a laboratory experiment to study the strategic use of vagueness in a voluntary disclosure game. Consider a sender who aims at inflating a receiver's estimate of her type and who may disclose any interval that contains her actual type. Theory predicts full information revelation if receivers are fully sophisticated. When facing a possibly naive receiver, the sender discloses an interval that separates her from worse types but is upwardly vague.

Senders in the experiment adopt this strategy and some (naive) receivers are systematically misled. Imposing precise disclosure in an additional treatment condition leads to less, but more easily interpretable, disclosure. Both theory and experimental data further suggest that imposing precision improves overall information transmission and is especially beneficial to naive receivers. Our results have implications for the rules that govern the disclosure of quality-relevant information by firms, the disclosure of research findings by scientists, and testimonies in a court of law.

While regulation is an important policy domain for which my approach delivers relevant insights, taxation is another important field of application. Taxation and tax compliance are essential for modern societies, yet there is much scope for improvement (Luttmer and Singhal, 2014). **Chapter 4**, which is based on joint work with *Charles Bellemare* and *Florian Englmaier*, investigates individual compliance behavior when facing complex rules, for example, when filing income tax returns. We focus on the question whether complexity contributes to self-serving compliance behavior like the underclaiming of benefits.

To study this question, we setup a laboratory experiment in which subjects face a compliance decision for which they need to file abstract tax forms of the Canadian province of Québec. The forms determine what share of a generated income can be kept by the subjects versus needs to be handed back to the experimenter (the "taxes" in our setting). In a fully factorial between subject design, we vary the complexity of these forms and for which purpose the collected taxes are to be used (i.e., the appropriation of tax money). This allows us to study whether the effects of complexity are linked to self/other-serving purposes, holding material compliance incentives constant.

Our results show a strong effect of the appropriation of taxes. Subjects are substantially less likely to comply with the tax rules when taxes are donated to a luxury yacht club rather than to a deserving cancer charity. We also observe that complexity causes a slight decrease in compliance. Interestingly, this decrease is particularly evident when taxes are distributed to the yacht club rather than to the cancer charity.

There exist at least two relevant explanations for this interaction effect. First, complexity might serve to alleviate the psychological costs from non-compliance. Second, the usage of taxes might change the motivation of subjects to comply with rules and hence they are more sensitive to complexity variations. When analyzing filing efforts and mistakes, we find suggestive evidence that is indicative of the first explanation. But in either case, the interaction effect of complexity and appropriation on compliance behavior has relevant implications for governments who jointly design tax rules and redistribute tax money.

Chapter 1

Cooperation in a Company: A Large-Scale Experiment[*]

1.1 Introduction

Within organizations most processes and production steps entail voluntary cooperation among employees to realize optimal output. This is particularly true for teamwork, but also for other daily interactions like helping or knowledge sharing (see Gittell, 2000; Fehr, 2018), where cooperation requires solving a social dilemma: those involved are better off if everybody provided high levels of effort or lots of time, but due to the individual incentive to contribute the enforceable minimum, the joint product is provided on a suboptimal scale, or not at all.

Social dilemmas have been studied extensively in the experimental laboratory (for reviews see Ledyard, 1995; Chaudhuri, 2011; Fehr and Schurtenberger, 2018) as well as in the context of governing the commons in the field (e.g., Ostrom, 1990; Rustagi et al., 2010; Fehr and Leibbrandt, 2011; Gneezy et al., 2016). Interestingly, there is much less empirical evidence on cooperation within organizations, and, in particular, companies.[1] They often have to solve a general tradeoff between creating a cooperative culture in order to provide internal public goods on an efficient level and securing a competitive environment in order to induce innovation and to be able to select the best employees for promotion. Striking the balance, given the tension between cooperation and competition, is probably one of the most difficult management tasks (Fehr and Fischbacher, 2002).

A key aspect of cooperation within organizations is that employees and teams often inter-

[*]This chapter is based on joint work with Martin Kocher and Christiane Schwieren.

[1]Notable exceptions are Charness and Villeval (2009) and Burks et al. (2009).

act repeatedly. While reputation concerns and informal peer sanctioning can reduce the free-rider problem, they are often unable to solve social dilemmas fully (e.g., Fischbacher and Gächter, 2010).[2] Thus, even in repeated interaction and with peer sanctioning mechanisms in place, it is essential for companies to establish a cooperative culture in order to sustain high levels of cooperation over time, avoiding the often observed decay in cooperation.

In this chapter, we exploit a unique setting for studying how financial and non-financial reward instruments within organizations relate to the cooperative culture among employees. Understanding this relationship entails relevant implications for many organizations. Our analysis is based on incentivized online experiments with 910 employees of a large software company.[3] We link data on the level of the employee from these experiments that measure cooperative attitudes in variants of the public goods game (see Fischbacher et al., 2001; Fischbacher and Gächter, 2010) with reward and context variables from company records.

Our setup allows for three main contributions. Firstly, we can systematically provide evidence on the association between cooperative attitudes and financial rewards within the company, while being able to control for determinants of cooperation whose relevance is suggested by economic theory. Secondly, we can assess potential non-financial reasons for cooperation in a natural environment that have so far almost exclusively been studied in the experimental laboratory. Thirdly, our study fulfills a methodological purpose by assessing the external validity for a business context of one of the most frequently applied laboratory measures of cooperation.[4]

With respect to our first contribution, we find that cooperativeness of employees does not lead to higher individual financial rewards. In stark contrast, our estimates show that within our study period from 2016 to 2018, cooperative employees received on average 29% lower annual wage increases and 15% lower financial award payments than their more selfish colleagues. Being cooperative is not rewarded but rather punished in terms of remuneration.

Regarding our second contribution, we observe that a large fraction of employees ex-

[2] Among other reasons, decreasing cooperation levels in repeated interaction result from contractual incompleteness of cooperative behavior (e.g., Holmstrom, 1982; Itoh, 1991), the existence of imperfectly conditional cooperators (see Fischbacher and Gächter, 2010; Ambrus and Pathak, 2011) or imperfect sanctioning mechanisms (Kandel and Lazear, 1992; Fehr and Rockenbach, 2003; Houser et al., 2008; Nikiforakis, 2008).

[3] Following the typology of Harrison and List (2004) our experiments can be referred to as an "artefactual field experiment". Alternatively, one could call it a "lab-in-the-field experiment" (Gneezy and Imas, 2017).

[4] There is an active methodological discussion about the generalizability/external validity of standard laboratory measures (see Levitt and List, 2007; Falk and Heckman, 2009; Burks et al., 2016; Gneezy and Imas, 2017).

hibits comparatively high levels of cooperation, despite the financial disincentives and the existence of selfish employees. Hence, in contrast to laboratory experiments, in which opportunistic cooperation is usually observed by selfish players in repeated cooperation that leads to a quick decay of contributions over time, we observe a potentially stable pattern of cooperation in the company. Consequently, behavior in the field experiment and observational data form the company together suggest that there must be substantial non-financial rewards of cooperation for the cooperators. Otherwise, cooperation should break down over time. While our online experiment features a one-shot interaction and thus cannot observe contribution dynamics, the high share of perfect conditional cooperators and the substantial number of unconditional cooperators provide the basis for stable cooperation.

We find supportive evidence for this interpretation when linking experimental data with record data from a non-financial recognition tool that employees can access via the company's intranet. Cooperative employees receive 51% more recognition awards from their colleagues. In a similar vein, we find that cooperative employees and teams comprised of a larger share of cooperative employees report stronger team cohesion and higher work satisfaction in our post-experimental survey, which is again a sign for non-financial reward components of a cooperative environment.

Regarding our third contribution, we document that cooperative employees send more than twice as many recognition awards than selfish employees. This correlation corroborates the external validity of cooperative attitudes measured in our experiments as sending an award requires some individual cost to write a justification and induces a positive externality on a co-worker.

Overall, our data is indicative of the idea that the company positively affects levels of cooperation through supplying non-financial compensating differentials to cooperative employees.[5] This is our preferred interpretation of the data, because it provides a joint mechanism for (i) high levels of cooperation, (ii) a negative nexus between financial rewards and cooperativeness, and (iii) a positive nexus between non-financial rewards and cooperativeness. We also investigate three other mechanisms that are likely to be present in our setting, but that are unlikely to be the sole driver of our three findings: an omitted variable bias related to performance or skills that are specific to cooperative attitudes (Bowles et al., 2001; Barr et al., 2009; Leibbrandt, 2012), selection based on cooperative attitudes (Falk and Heckman, 2009; Dohmen and Falk, 2011), and context-dependent pref-

[5]This interpretation is in line with a strand of literature that emphasizes an intrinsic value of cooperation beyond its financial consequences (Hamilton et al., 2003; Bandiera et al., 2005, 2011, 2013; Ruff and Fehr, 2014).

erences (Bowles, 1998; Levitt and List, 2007; Cohn et al., 2014).

The remainder of this chapter is structured as follows. We first relate our study to the literature on artefactual field experiments to study cooperation in the field. In Section 1.3, we outline the company setting at hand. In Section 1.4, we describe our experimental setup and the data for our analysis. Then, we report the correlation between cooperative attitudes and relevant outcome variables from the company context in Section 1.5. Section 1.6 discusses the main findings and potential underlying mechanisms. Section 1.7 concludes the chapter.

1.2 Related Literature

It is impossible to do justice to the large experimental literature on cooperation, even if one restricts attention to (artefactual) field experiments and lab experiments predicting prosocial behavior outside the laboratory (for a survey see Galizzi and Navarro-Martínez, 2019). Examples for field experiments on cooperation are List and Lucking-Reiley (2002), Cardenas (2003), Frey and Meier (2004), Alpizar et al. (2008), Benz and Meier (2008), Burks et al. (2009), Carpenter and Seki (2011), Croson and Shang (2008), Rustagi et al. (2010), Fehr and Leibbrandt (2011), Voors et al. (2011, 2012), Stoop et al. (2012), or Gneezy et al. (2014, 2016). People have studied charitable giving, fishermen, truck drivers, visitors of national parks and many more. However, there is very little evidence on company settings.

Regarding our main research interest, the financial and non-financial rewards of cooperative attitudes of employees in a company there is particularly scarce existing empirical evidence from the field. This is despite an abundance of case studies and anecdotical evidence on firms that must balance cooperative and competitive elements in their incentive schemes or that must foster cooperation within teams to be successful (e.g., Dirks, 1999; Dirks and Ferrin, 2001; Gratton, 2009, 2011; Grant, 2013). Beersma et al. (2003) discuss the relevant management literature and provide a study on the cooperation/competition tradeoff, including personality differences and task characteristics.

In the following, we provide an upshot of the existing literature on our three main contributions. Our first contribution is on the association between cooperative attitudes and financial rewards within the company. Burks et al. (2009) use a naturally occurring social dilemma among bicycle messengers in Switzerland and the United States. Their focus is on the selection of messengers into companies based on incentive schemes. Workers in companies that pay for performance show less cooperation than workers in companies that pay fixed hourly wages or that are members of cooperatives.

There is more closely related literature in other than a standard workplace domain (or using other paradigms than the standard public goods game) that can still inform our setup. Leibbrandt (2012) compares behavior of professional shrimp sellers in a laboratory public goods game with natural market outcomes. He finds a positive relationship between cooperativeness and market success as measured by achieving higher prices for shrimps and establishing longer lasting trade relations. He argues that the detected correlation is driven by cooperative employees being able to signal trustworthiness. Similarly, Essl et al. (2018) study the trustworthiness of sales employees of an Austrian retail chain using a modified trust game and relate behavior in the game to individual sales performance data. The authors find that higher trustworthiness is associated with lower sales per day, but with higher revenue per customer. Cardenas and Carpenter (2005) look at experimental measures of cooperation and link them to household expenditures in Vietnam and Thailand, showing that more cooperative individuals are better off. Likewise, Barr and Serneels (2009) provide evidence that experimentally elicited trustworthiness is positively related to wages of manufacturing workers in Ghana.

Regarding the relationship between cooperative attitudes and non-financial rewards – our second contribution – there again exists limited evidence. Ruff and Fehr (2014) summarize evidence from laboratory FMRI studies that indicate "[...] an experienced value of cooperation per se that might bias individuals to display cooperative behavior" (p. 557). In the field, Hamilton et al. (2003) show that workers at a garment plant voluntarily select into a team-based work organization despite financial losses as compared to performing sewing tasks individually. They argue that such selection behavior is likely driven by non-financial reasons such as hedonic benefits from team work. In a similar vein, Bandiera et al. (2005, 2011, 2013) find that UK fruit pickers increase efforts or forgo financial benefits due to social ties to co-workers.

Our third contribution relates to the external validity of experimentally elicited cooperative attitudes. While we know quite a lot on the external validity of different measures on uncertainty preferences (risk and ambiguity) and time preferences, we know much less on the external validity of standard measures of cooperative attitudes. Existing studies that provide evidence of the external validity of the standard linear public goods game, i.e., the voluntary contribution mechanism (VCM), are mainly linked to the problem of the commons. Rustagi et al. (2010) elicit cooperative attitudes of members of 49 forest user groups in Ethiopia in an artefactual field experiment setting. They link cooperative attitudes to natural forest commons outcomes and find that groups that are comprised of a larger number of conditional cooperators are doing a better job in managing the forest commons. In a similar vein, Gneezy et al. (2016) study Brazilian fishermen who are organized differently in different places regarding the need for team work. Fishermen at the sea who are forced to work in teams cooperate and trust more than their counterparts at lakes who mostly work individually (see also, for instance, Carpenter and Seki, 2011; Fehr and Leibbrandt, 2011; Stoop et al., 2012; Voors et al., 2011, 2012).

Evidence for contexts, apart from common pool management, is provided by Burks et al. (2016) who conduct prisoner's dilemma experiments with truck drivers. More cooperative truck drivers are found to send satellite uplink messages from their trucks more frequently (messages are costly but benefit an anonymous colleague). Englmaier and Gebhardt (2016) perform a lab-field comparison by inviting student participants to a laboratory public goods game and to a natural work setting (registering books in a library database) in which incentives condition on team outcomes. From the positive correlation of behavior in the laboratory and in the natural work task, the authors conclude that the laboratory public goods game captures important aspects of structurally equivalent situations outside the laboratory.[6]

Our study adds the novel elements of a work team and company setting and the link of financial as well as non-financial outcome data with behavioral data from incentivized experiments to the existing literature in economics and management.

1.3 The Company Setting

We conduct our study in partnership with a large, multinational software company. About 40% of the employees work as software developers, 40% work in the sales and consulting area, and the remaining 20% work in more general service areas like Human Resources, Accounting & Finance or Marketing. Several institutional features are important to understand how the company and its reward systems operate.

Business Models. Most individual and teamwork tasks in the company are mainly taking place in either a customer business model or cloud business model. The customer model uses servers that are on the premise of the client and that are serviced by company employees, whereas the cloud business model uses internet cloud solutions that concurrently apply to many clients. According to our discussions with managers of the company before conducting the study, the latter model requires more cooperation among workers at the software producer than the former; in other words, it entails a production function with much more pronounced complementarities (for instance, between software

[6] As in our study, Charness and Villeval (2009) deploy a linear public goods game in actual companies, but they focus on the difference in behavior of junior and senior employees. The main finding is that senior employees are more cooperative than junior employees. Von Bieberstein et al. (2020) analyze student performance in math exams and partner work assignments at university using public goods game measures, however, they find no correlation (but free-riders are performing better in the exam).

development and consulting).[7] Interestingly, due to the cloud model connecting several software products on an interface, sales employees also have sales bags that are comprised of items that, if sold, positively affect the performance of their sales team, i.e. other team members.

Pay Schemes. Employees are enrolled in one of two co-existing pay schemes: either the company performance or the individual performance pay scheme. Both schemes involve a fixed component and a variable pay component. They differ in how the variable pay component is determined. In the company performance pay scheme, employees receive bonus payments that are determined by the overall company performance. Under the individual performance pay, bonuses depend on individual performance assessments. Enrollment in either of these schemes is tied to job roles such that selection is only possible via job choice. While all developers and employees in the service areas work under the company performance pay, most sales employees work under individual performance pay. Consultants are equally likely working in either of the schemes depending on whether they are in-house or outgoing consultants.

Wages. The employees' target wage consists of a base wage and the bonus conditional on full target achievement (either company or individual target). This means that the target wage does not necessarily correspond to the actual wage payed. However, analyses by the company show that the target wage is a good proxy for actual wages, hence, we refer to them as wages.[8] Cross-sectional variation in wages is mainly due to jobs at different career levels or in different functions. Variations in the wage levels over time reflect job trajectories. For example, this includes promotions or other internal job changes that relate to a different pay mix. In addition, managers have a budget for merit increases paid to their employees to be decided upon on a yearly basis.

Financial Awards. Another important reward instrument of managers is the conferral of financial awards. At the end of a year, every manager can allocate financial awards that consist of shares of the company among employees in his/her team. An award conferred in a particular year is paid out in three tranches in the subsequent years. The budget is fixed for each year for the whole company and on team levels. The award guidelines

[7] For validation, we ask all participants how important cooperation is to successfully fulfill their individual and teamwork tasks on a standard Likert scale in an online survey. We detect a strong correlation between the business models and responses to the survey question (Spearman correlation: $-0.214, p < 0.001$). While 42% of employees state that teamwork is of high importance in the cloud business model only 24% do so in the customer business model (t-test, $p < 0.001$).

[8] In the company performance scheme, there was full target achievement over the relevant years; hence, target wages equal the wages payed. In the individual performance pay, target wage is a noisier measure of the actual wage payed. While on average there are very high target achievement rates (on average, above 100%), there is a higher standard deviation.

handed out to the managers specify the idea of a financial award as recognizing employees that are important for the success of the company and as an instrument for employee retention. The guidelines apply to all departments, job positions and both pay schemes.

Recognition Awards. Furthermore, there exists a non-financial recognition system that every employee can easily access via the company's intranet. The program is an institutionalized way to thank a colleague for several desired behaviors including, for example, cooperation, promise keeping, or embracing diversity. If an employee receives an award, he/she is notified via e-mail. The e-mail prominently shows a slogan such as "Thank you for being cooperative!" (or the relevant other award justification). It also contains a message from the sending employee and his/her name. The receiving employee's manager can see every award and the total number of awards received for each team member. The role of the manager is also to prevent employees from sending awards back and forth. There are no direct financial consequences related to a recognition award, neither for the sending nor for the receiving employee. However, sending an award requires some effort as it must be justified in a text of at least 150 characters.

1.4 Experimental Setup and Data

Our analyses are based on data from three different sources. First, we collect data from an incentivized online experiment. Second, in a subsequent survey module, we elicit a variety of control variables such as socio-demographic characteristics or behavioral measures that relate to cooperation. The gathered data is then merged with reward and context variables from the company records on the individual level. An overview of all collected variables can be found in Appendix A.1.

1.4.1 Behavioral Measure of Cooperative Attitudes

The first part is a public goods experiment according to the "ABC-framework of cooperation" (Gächter et al., 2017).[9] It uses the design of Fischbacher et al. (2001), including the elicitation of beliefs. In a VCM setting, we elicit an unconditional contribution to a public good, a full contribution schedule contingent on average contributions of other group members, and subjects' beliefs about others' average unconditional contributions.

Participants are randomly assigned to groups of three. Every participant knows that all other participants are randomly selected employees of the company. Each group member

[9]The instructions of the public goods game can be found in Appendix B.2. The full experimental material provided to employees can be found in the Online Appendix.

receives an initial endowment of 10 Tokens to be allocated to a private account or to be contributed to a public account. One Token equals €1. The invested amount $c_i \in \{0, 1, ..., 10\}$ is referred to as the unconditional contribution. The sum of all contributions to the public good is multiplied 1.5 in our case, and divided equally among all $n = 3$ group members. This leads to the following payoff function for subject i:

$$\pi_i = 10 - c_i + \gamma \sum_{j=1}^{n} c_j \tag{1.1}$$

which is linear in the public good contribution and where c_i denotes the contribution of group member i. The marginal per capita return (MPCR) from investing in the public good is $1/n < \gamma = 0.5 < 1$. From an individual perspective, free-riding (i.e., $c_i = 0$) is a dominant strategy. Since the sum of marginal returns is larger than 1, however, contributing the entire endowment (i.e., $c_i = 10$) is the optimal choice from a collective perspective. The decision is made only once and anonymously. Thus, there are no incentives and no possibilities to build a reputation.

Participants do not receive any feedback after indicating an unconditional contribution. Subsequently, participants are asked to fill in a contribution table indicating their contribution for each possible average contribution of the other group members, rounded up to integers. The conditional contributions from the contribution table allow us to classify three distinct cooperative attitudes. We depart from the existing literature for expositional reasons; the interpretation of our analysis is simplified when using the three categories. Fischbacher et al. (2001) and many follow-up papers classify free riders (zero contributions, regardless of the average contributions of others), conditional cooperators (increasing contributions with increasing average contributions of others), and hump-shaped contributors (increasing contributions with increasing average contributions of others up to a certain contribution level, and above decreasing contributions with increasing average contributions of others). Since we additionally observe a significant number of perfect conditional contributors (those who match the average contributions of others perfectly) and even some unconditional full contributors (contributing the maximum amount of ten Tokens regardless of the average contribution of others), we use the following classification:

- Net-Taker: We classify an employee whose average conditional contribution is smaller than five Tokens as a Net-Taker. This means that the employee, on average, free-rides (at least partially) on the contribution of others to the public good (mainly free riders, conditional cooperators with a self-serving bias[10]).

- Net-Giver: An employee that contributes more than five Tokens is defined as a Net-Giver. The employee, on average, contributes more than the two others (mainly conditional cooperators with an other-serving bias, unconditional full contributors).

- Matcher: An employee who, on average, exactly matches the average contribution of the two other members is considered a Matcher (almost equivalent to perfect conditional cooperators).[11]

To make both the unconditional and the conditional contributions incentive-compatible, we use the mechanisms described in Fischbacher et al. (2001). That is, for one randomly selected subject the conditional contributions are payoff-relevant, whereas for the two remaining subjects the unconditional contribution is to determine the average contribution of other group members. We also elicit expected contributions of others in an incentivized way. Following Gächter and Renner (2010), participants are asked to guess the average unconditional contribution of the other group members and receive €5 if they are correct, otherwise they receive €0.

1.4.2 Survey and Company Variables

After the incentivized parts, we elicit additional variables that are relevant for the analysis of the determinants and the context of cooperation without using monetary incentives. Importantly, we ask whether an employee's individual and teamwork tasks are mainly related to the customer or the cloud business model. In addition, we capture personality traits (a short form of the Big Five; Rammstedt et al. (2013)), and survey measures of related social preference concepts like negative/positive reciprocity (Falk et al., 2018) and trust (Anderson et al., 2004). We also elicit a measure of individual competitive attitude (i.e., the competitiveness index as introduced by Newby and Klein (2014)) and basic socio-demographic variables (such as nationality, education, and number of kids and friends). Furthermore, variables with respect to perceived team cohesion (Ashforth and Mael, 1989), team stability, and work-related stress (Schulz and Schlotz, 1999) are elicited.[12]

[10]These are conditional cooperators that have an increasing contribution schedule, but they, on average, contribute less than the average of other members.

[11]Our main results are robust to using different definitions of cooperative attitudes. Details are provided in the results section and in the Appendix.

[12]After the main public goods game, we also use incentivized coordination games and short social dilemma vignettes to elicit the shared perception of cooperative norms that prevail in the company (compare to

We combine the elicited data with a rich data set from the company. On the employee level, this includes age, gender, seniority (years employed at the company), career levels, and personal leadership responsibilities. Using a work team identifier, we can also infer information about team compositions (for example, with respect to gender and age). Regarding reward institutions, we have individual level information on the employees' pay scheme, his/her wage level, and the value of financial award payments. Observing employees' wage levels over time allows us to calculate annual wage increases. We additionally observe the numbers of recognition awards received and sent for each employee.

1.4.3 Procedures

We conducted the described experimental and survey modules online.[13] Eligible employees received a personalized participation link. Every respondent knew that he/she can complete the experiment within a two-weeks period. There were two roll-out phases with different employees, the first in November 2017 and the second in February 2019. Employees could participate during regular work hours. The total completion of the experiment and the survey took about 30 minutes and could be interrupted at any time.

The online experiment did not require participants to simultaneously make decisions. Participants were informed that groups were assembled randomly ex post. Since nobody received feedback during the experiment, such a procedure is equivalent to simultaneously entered decisions. Participants could use their personal ID code to login after the roll-out phase had ended to get feedback on the results. We asked participants to perform the online experiment individually. The random and anonymous allocation to groups made sure that coalition formation among group members when filling in the online experiment was impossible.[14]

Before a participant could decide about the public good contributions, he/she needed to answer comprehension questions on the game. If an answer was wrong, the participant was notified and was shown the correct answer to be re-entered in the respective input box. We set up a telephone hotline and an e-mail address for potential questions during

Burks and Krupka, 2012). This provides us with a better understanding of the "cooperative culture" in the company. For an extensive discussion of these norm elicitations and the respective empirical results refer to Deversi et al. (2020b).

[13]Our study represents one of many studies and surveys that employees fill out at the company. The company even has its' own survey team. Hence, asking employees to participate in an online study while being at their workplaces is nothing unusual, although the incentivized experimental part was of course somewhat special to most employees.

[14]It was extremely unlikely that (matched) participants would be sitting in a shared office. Analyses of the participants' start and end times suggest that there was no communication or coordination of employees of a work team (for the analysis see Appendix A.3).

the experiment. We received very few calls and messages.

In the first roll-out phase in 2017, we implemented an unexpected donation option at the end of the experiment as a control for social desirability concerns. In 2019, we included an additional public goods game (administered in a within-subject fashion) that varied the MPCRs (either very high, 1.2, or very low, 0.3) to check whether participants would react to changes in the the social dilemma characteristics. Notice that an MPCR of 1.2 makes it individually optimal to contribute, whereas an MPCR of 0.3 makes it both individually and social optimal not to contribute.

Individual data from the company was de-identified before linking it to our elicited data. The data collection and storage were facilitated through Qualtrics. There exists a data protection agreement between the company and Qualtrics; and a research agreement (including data protection) between the company and the research team. Data protection units at the company, at University of Munich and University of Heidelberg supervised the study. The company did not receive individual-level data, and all participants were informed about the full pseudonymization of their responses before the experiment. The data protection at the company was only to be involved in determining the exact procedures, not in handling the linked data. We made sure that the pseudonymized final data set was only stored on the computers of the researchers involved in this project within university fire-walls.

Employees were aware of the data protection procedures and provided informed consent before participating in the study. Ethics approval by the University of Munich was granted in September 2017. The study was pre-registered at the AEA registry (AEARCTR-0002596). The respective pre-analysis plan was slightly updated and re-submitted before the second round of experiments took place in 2019.

1.4.4 Sample and Selection

We invited 2,799 employees from 371 work teams to participate in our study.[15] This includes 1,297 employees invited in 2017 and 1,502 employees invited in 2019. We randomly selected teams that had between 8-20 team members of which more than 70% were based in the German-speaking area.

[15] In 2019, we excluded working students and temporarily employed consultants from invitations. Also, in 2017, we slightly oversampled employees from the individual performance pay scheme to have a larger comparison group. There was limited record data availability for these employees in 2017. Working students and external employees were not eligible to participate in award programs and worked under special fixed wage contracts. Hence, we decided to exclude these groups of employees from the second round of experiments in 2019.

Overall, 910 employees from 299 teams participated.[16] This corresponds to a participation rate of about 32.5%. The characteristics of the participating employees are mostly representative for the employee population at the company (conditional on the invitation requirements) as can be seen in Table 1.1. There does not seem to be any selection bias into the experiment based on observable characteristics. However, compared to non-participating employees, participating employees less frequently work under the individual performance pay scheme (26% versus 22%). Almost all participating employees are placed in the German-speaking area (99% versus 98% in the invited sample). We did not receive wage data for 57 participating employees. These data were either secret, from working students or external employees, or were not available to the company's German human resources department that we worked with to retrieve the data from the records.

More generally, one might expect sample selection according to the unobserved level of cooperativeness of employees. Cooperative employees could more frequently volunteer to participate in surveys/experiments, which could bias our results and interpretations. First, this is not so much of a concern, given that we are not interested in the level of cooperation, but in the link between cooperation and company outcomes. Second, as a robustness check, we show in Section 1.5.2 that given a high correlation between, for example, recognition award sending and cooperativeness, we do not find any evidence for the systematic selection into our experiments based on cooperative attitudes. The significant correlates of cooperativeness are statistically indistinguishable between participating and non-participating employees.

1.5 Results

1.5.1 Cooperative Attitudes

About 24% (N=201) of the employees can be classified as Net-Takers, i.e., they contribute on average less than five Tokens (mean of 2.51 Tokens) in the conditional contribution decisions. We classify 35% (N=345) as Net-Givers who contribute on average more than five Tokens (mean of 7.23 Tokens). Around 41% (N=364) of the employees exhibit a contribution pattern best described by Matcher behavior, which implies an average contribution of exactly five Tokens.[17]

[16] We count an employee's response as a full response if more than 90% of the questions were answered. Herewith we exclude 414 employees that answered on average only 9.8% of the questions – which corresponds to the first screen of the public goods game instructions.

[17] We observe similar distributions of cooperative attitudes comparing the experimental waves in 2017 and 2019 (Komoglorov-Smirnov Test, $p = 1.000$). This also holds for the other public goods game variables. Hence, for the period of our study, we regard the cooperation pattern in the company as stable and pool

	Non-Participating Employees			Participating Employees			
	Count	Mean	SD	Count	Mean	SD	P-Value
Socio-Demographics							
Female	1889	0.30	0.46	910	0.30	0.46	0.965
Age	1889	45.09	8.95	910	44.48	9.31	0.276
Company Controls							
Seniority	1889	14.33	7.34	910	14.03	7.47	0.292
Team Size	1889	13.60	3.54	910	13.78	3.48	0.264
Leader	1889	0.09	0.29	910	0.10	0.30	0.571
Career Level							
Low	231	0.12	0.33	111	0.12	0.33	0.993
Medium	1430	0.76	0.43	683	0.75	0.43	0.762
High	228	0.12	0.33	116	0.13	0.33	0.686
Indv. Performance Pay	1772	0.26	0.44	866	0.22	0.41	0.031
German Area	1889	0.98	0.01	910	0.99	0.01	0.041
Outcome Variables							
Recognition Awards							
Reception	1892	0.29	0.89	910	0.26	0.61	0.823
Sending	1892	0.21	1.51	910	0.22	1.19	0.623
Wage	1779	.	.	853	.	.	0.217
Wage Increase	1774	0.044	0.078	846	0.045	0.086	0.154
Financial Awards	1774	0.058	0.057	873	0.061	0.055	0.150
N	1889			910			

Table 1.1 Sample Selection

Notes: P-values rely on two-sample Mann-Whitney-U tests for continuous variables or χ^2-tests for categorical variables. For reasons of discretion, we do not provide wage level statistics here. However, there is no significant difference in wage levels between participants and non-participants. Career levels subsume several actual categories in each presented category. Financial awards are denominated in percent of wages.

Table 1.2 presents an overview of the collected public goods measures for each of the three cooperative types. Overall the unconditional contribution decisions reveal very high cooperation levels (79% of the endowment), despite the existence of Net-Takers. Net-Takers contribute significantly less unconditionally than Matchers and Net-Givers (5.44 versus 8.41 and 8.77, respectively; Mann-Whitney-U (MWU) tests, both p-values < 0.001). They also expect lower unconditional contributions from their colleagues (4.54 versus 7.30 and 7.32, respectively; MWU tests, both p-values < 0.001). Differences between Matchers and Net-Givers are not statistically significant (MWU tests; unconditional contributions, $p = 0.876$; beliefs, $p = 0.436$).

Following Fischbacher and Gächter (2010), we estimate each employee's slope parameter from a linear regression of the conditional contribution and the contribution schedule.[18]

the data whenever possible.

[18]If the slope parameter is equal to 1, all contributions of the employee coincide with the average contribution

	All (N=910)		Net-Takers (N=201)		Matchers (N=364)		Givers (N=345)	
	Mean	Sd	Mean	Sd	Mean	Sd	Mean	Sd
Unconditional contributions	7.89	2.93	5.44	3.54	8.41	2.58	8.77	1.96
Belief about others' contributions	6.70	2.78	4.54	2.79	7.30	2.57	7.32	2.34
Mean conditional contribution	5.30	2.25	2.51	1.76	5.00	0.00	7.23	1.77
Slope parameter	0.71	0.43	0.46	0.43	0.95	0.24	0.59	0.45

Table 1.2 Overview of Public Goods Game Variables by Cooperative Attitudes

The average slope parameter is 0.71, which reflects a tendency to conditioning own contributions on others' contributions. The Net-Takers' average slope parameter equals 0.46 and is lower than the parameters of the other two attitude types (MWU tests, both p-values < 0.001). While the Matchers' slope parameter is almost 1 (mean of 0.95), reflecting that most of these employees are perfectly conditionally cooperative, the Net-Givers have a slope parameter of 0.59, which lies between the other two attitude types (MWU tests, all p-values < 0.001).[19]

1.5.2 Recognition Awards and Cooperative Attitudes

Figure 1.1 relates the number of received (left) and sent (right) recognition awards per employee to cooperative attitudes. We observe that Net-Givers act more cooperatively and are also recognized as such. They sent more than 2.5 times as many recognition awards and receive about 40% more than their colleagues (MWU tests, pooling Net-Takers and Matchers, $p = 0.057$ and $p = 0.039$, respectively). The difference between Net-Givers and Matchers in sending behavior is statistically significant (MWU test, $p = 0.012$), and the difference between Net-Givers and Net-Takers in reception levels is as well (MWU test, $p = 0.053$).

of the other two group members, i.e., there is a perfect linear relationship between their contribution and the contributions of the others (perfect conditional cooperation). If the parameter decreases the relationship becomes weaker, such that a value of 0 means that contributions are independent of the others' average contribution.

[19]In Appendix A.4, we show in more detail how our cooperative attitudes are related to the cooperation types proposed by Fischbacher et al. (2001) and Fischbacher and Gächter (2010). In Appendix A.5, we provide an extensive multivariate analysis that characterizes cooperative attitudes in terms of the employees' personal, behavioral and work-related characteristics. We observe a positive relationship between age and cooperativeness, and interestingly that female employees are less cooperative than male employees. In terms of the behavioral survey measures, we document that employees are more likely to be Net-Takers the more competitive, distrusting, negatively reciprocal, extroverted and neurotic they are. Besides, we find that employees in the individual performance pay scheme are more likely Net-Takers than Matchers as compared to employees in the company performance scheme. There are no significant differences in the distribution of cooperative attitudes with respect to career levels, leadership responsibility, seniority, or business model.

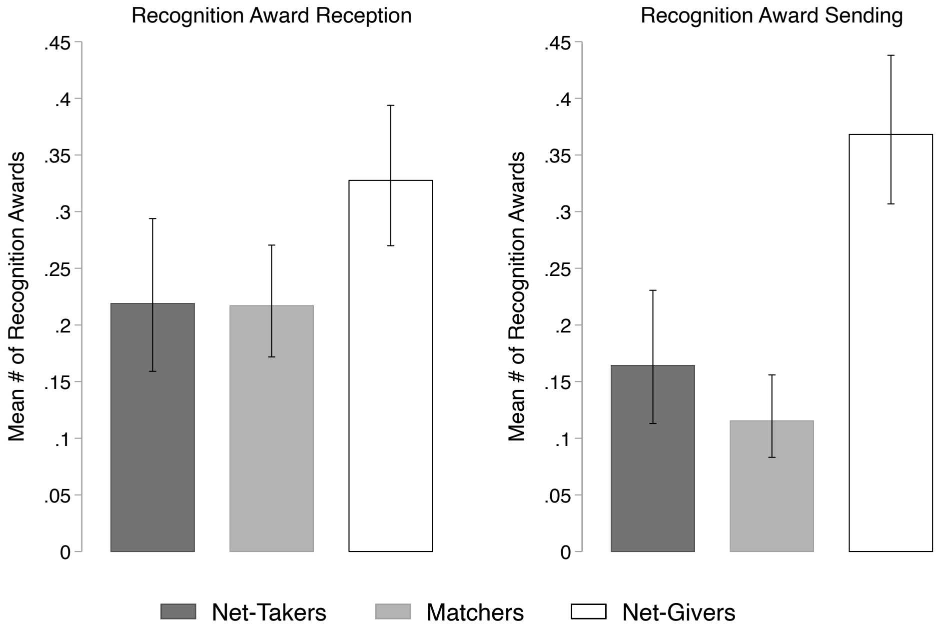

Figure 1.1 Recognition Awards and Cooperative Attitudes

Notes: The graph bar contains data on recognition awards from 2017 for participants in the experiments in 2017 and data from 2018 for participants in 2019. Bars show sample means for each cooperative attitude. Vertical caps show the 95%-confidence interval that is calculated based on a Poisson distribution.

We model the number of received (R_r) and sent (R_s) recognition awards as

$$E(R_{(r,i)}|\boldsymbol{X}_i) = \exp\left(\alpha + \beta'_{(r,1)}\boldsymbol{C}_i + \beta'_{(r,2)}\boldsymbol{X}_i + \beta_3 \text{year}_i\right) \tag{1.2}$$

$$E(R_{(s,i)}|\boldsymbol{X}_i) = \exp\left(\alpha + \beta'_{(s,1)}\boldsymbol{C}_i + \beta'_{(s,2)}\boldsymbol{X}_i + \beta_3 \text{year}_i\right) \tag{1.3}$$

where $\boldsymbol{C}$ is the vector of dummies for Matchers and Net-Givers using Net-Takers as the base category. The covariate vector $\boldsymbol{X}$ consists of socio-demographics and company controls, including the career level and job role as defined by the department (e.g. software development). The variable year absorbs differences between 2017 and 2018.

The respective multivariate Poisson regression estimations presented in Table 1.3 are in line with the preceding non-parametric analyses. Net-Givers receive 51% more awards and send more than twice as many awards as Net-Takers, when including socio demographics and company controls (see columns (3) and (6)). Due to relatively low number of employees sending awards (about 11% sent at least one award), these estimates are less precise then the estimations for the reception patterns. Notably, we also observe that Matchers receive and sent significantly fewer awards than Net-Givers.

We take the comparatively high number of sent recognition awards by Net-Givers as ev-

	(1) # Awards Received	(2) # Awards Received	(3) # Awards Received	(4) # Awards Sent	(5) # Awards Sent	(6) # Awards Sent
Net-Taker	0	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)	(.)
Matchers	-0.012	-0.009	-0.060	-0.358	-0.377	-0.445
	(0.214)	(0.214)	(0.224)	(0.405)	(0.407)	(0.388)
Net-Giver	0.403**	0.429**	0.413**	0.807*	0.789*	0.807**
	(0.191)	(0.194)	(0.195)	(0.431)	(0.433)	(0.379)
Constant	-1.829***	-1.892***	24.40**	-2.332***	-3.207***	28.89**
	(0.363)	(0.504)	(10.32)	(0.657)	(0.878)	(13.74)
b[Matchers] - b[Net-Givers]	p=0.016	p=0.013	p=0.013	p=0.001	p=0.001	p=0.002
Socio Demographics	X	✓	✓	X	✓	✓
Company Controls	X	X	✓**	X	X	✓**
Career Dummies	X	X	✓***	X	X	✓***
Dep. Dummies	X	X	✓***	X	X	✓***
Observations	910	907	842	910	907	842
Pseudo R2	0.042	0.048	0.080	0.056	0.063	0.157
Model	Poisson	Poisson	Poisson	Poisson	Poisson	Poisson

Table 1.3 Regressions of Recognition Awards on Cooperative Attitudes

Notes: Standard errors clustered on team-level in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; Asterisks for the control variables show the test result from an F-Test, testing the joint difference from zero. Alternative estimations using zero-inflated poisson models yield qualitatively very similar results.

idence for the external validity of experimentally elicited cooperation levels. Sending an award induces a positive externality on a co-worker and requires writing a justification for the award, i.e., it represents a costly pro-social act similar to public goods game contributions. The externality may involve positive emotions on the recipient's side, but also potentially some indirect monetary value. Remember that managers observe awards; hence, monetary consequences could include financial awards and merit increases, respectively. Moreover, recognition awards seem to be unrelated to a strong reciprocity concern as Matchers, who exhibit strong reciprocity through their contribution schedule, receive and send significantly fewer awards than Net-Givers.

1.5.3 Financial Rewards and Cooperative Attitudes

Figure 1.2 shows the mean annual wage increases and the financial award allocation by cooperative attitudes and pay schemes. A similar pattern arises for both variables:[20] When pooling data from both pay schemes, Net-Takers receive a higher financial appre-

[20] The Spearman correlation coefficient between wage increases and financial awards is rather weak at 0.081, but still statistically different from zero at the 5% level ($p = 0.019$).

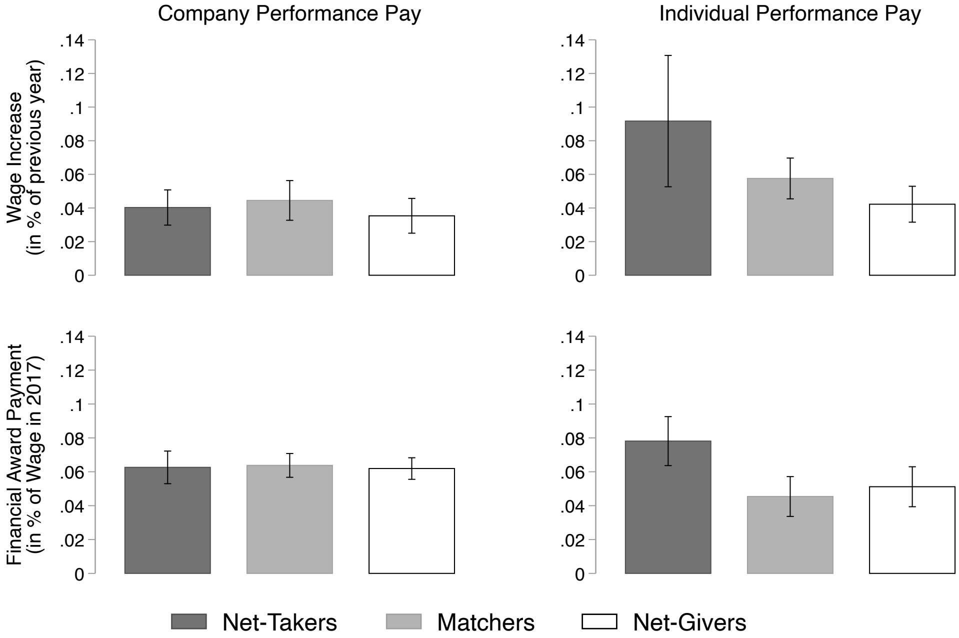

Figure 1.2 Financial Rewards and Cooperative Attitudes

Notes: Bars show sample means for each cooperative attitude. Vertical caps show the 95%-confidence interval that is calculated based on a standard normal distribution. (Top) The graph bars contains data from 2016-2018 for participating employees. (Bottom) The graph bars contain data from 2017 for all participating employees.

ciation than their colleagues (MWU tests, wage increases, $p < 0.001$; financial awards, $p = 0.077$). Focusing only on company performance pay, we detect no heterogeneity with respect to cooperative attitudes (MWU tests, lowest p-value$= 0.534$). Focusing only on individual performance pay, Net-Takers receive significantly higher financial rewards than other employees (MWU tests, pooled, for both outcomes $p < 0.001$). This also holds when comparing Net-Takers with Net-Givers (MWU tests, wage increase, $p < 0.001$; financial awards, $p = 0.009$) and Matchers separately (MWU tests, wage increase, $p = 0.037$; financial awards, $p < 0.001$).

We model the financial appreciation variables using linear regressions. Wage increases ($\frac{w_t}{w_{t-1}}$) are measured in percent of the base year (either 2016 or 2017 depending on the year of participation). Financial award payments (f) are measured in percent of the wage in 2017.

$$\left(\frac{w_t}{w_{t-1}}\right) = \alpha + \beta'_{(1}\boldsymbol{C}_i + \beta'_2\boldsymbol{X}_i + \beta_3\text{year}_i + \varepsilon_i \tag{1.4}$$

$$\left(\frac{f}{w_{2017}}\right) = \alpha + \beta'_1\boldsymbol{C}_i + \beta'_2\boldsymbol{X}_i + \varepsilon_i \tag{1.5}$$

In model (4), we use the same covariates as described in model (2).[21] In model (5), we drop year dummies as we include data on financial award payments from 2017 only.

Table 1.4 shows estimated coefficients from OLS regression models. Columns (2) and (6) contain estimated differences between cooperative attitudes, while controlling for socio-demographic and company covariates. We observe that Net-Givers' wage increases are 29% (1.5%-points) and financial award payments are 15% (1%-point) lower than Net-Takers' appreciation, respectively. As already suggested by Figure 1.2, this difference is only relevant in the individual performance pay scheme. Here, Net-Givers receive about 48% (4.4%-points) lower wage increases and 32% (2.7%-points) lower financial award payments than Net-Takers (see columns (4) and (8), respectively). We observe no differential financial appreciation between Matchers and Net-Givers and no differences in the company performance pay scheme.

One can also look at whether cooperative attitudes and observables determine *wage levels* instead of *wage increases*. The results of the analysis are provided in Appendix A.8. Controlling for relevant Career and Department Dummies as well as socio-demographic and company control variables, only age is a significant determinant of overall wage levels. There is no significant interaction effect with the incentive scheme, either. Obviously, short-term changes in the wage levels are much more responsive to cooperative attitudes. We know that these variations might change with age, with incentive schemes, and with other influences. Together with potential long-term selection effects into different areas or jobs within and outside the company and leveling effects of collective bargaining agreements over time that matter for the overall wage levels, regressions that use wage levels as dependent variable are probably not that informative for our setup. Hence, the results based on wage levels should be interpreted carefully; we would have needed a much more flexible wage determination environment (e.g., top-level management) to detect a potential relationship between cooperative attitudes and wage levels.

1.6 Analysis of Potential Mechanisms

How can a company achieve high levels of cooperation despite financial disincentives to cooperate? According to Rosen (1986), teamwork at the workplace (and cooperation) involves other, non-financial returns for employees such as less boring work or hedonic benefits from social interaction. In the context of our study, such non-financial returns (e.g. measured by the number of received recognition awards) are likely to act as equaliz-

[21] In Appendix A.7, we include the change in part-time shares for years 2016/2017 and 2017/2018 as a covariate for wage increases. This allows us to control for employees that moved to parental leave or partial retirement during the period of our study. The results remain largely robust.

	(1) Wage Increase (in %)	(2) Wage Increase (in %)	(3) Wage Increase (in %)	(4) Wage Increase (in %)	(5) Fin. Award Payment (in % of Wage)	(6) Fin. Award Payment (in % of Wage)	(7) Fin. Award Payment (in % of Wage)	(8) Fin. Award Payment (in % of Wage)
Net-Takers	0	0	0	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Matchers	-0.00820	-0.00475	0.00465	0.00533	-0.00666	-0.00702	0.00134	0.000536
	(0.00859)	(0.00791)	(0.00831)	(0.00792)	(0.00526)	(0.00435)	(0.00636)	(0.00509)
Net-Givers	-0.0190**	-0.0150**	-0.00439	-0.00344	-0.00719	-0.0102**	-0.000422	-0.00313
	(0.00776)	(0.00744)	(0.00716)	(0.00757)	(0.00486)	(0.00475)	(0.00577)	(0.00547)
Ind. Perf. Pay		0.0228**	0.00896	0.00896		0.00289	-0.00922	-0.00315
		(0.00980)	(0.00831)	(0.0105)		(0.00662)	(0.00749)	(0.00823)
Net-Takers × Ind. Perf. Pay			0.0454**	0.0442**			0.0268**	0.0269***
			(0.0203)	(0.0197)			(0.0105)	(0.0103)
Matchers × Ind. Perf. Pay			0.00620	0.00561			-0.00754	-0.00347
			(0.0119)	(0.0115)			(0.00955)	(0.00928)
Net-Givers × Ind. Perf. Pay	0	0	0	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Constant	0.0536***	0.0520**	0.0287**	0.0421	0.0576***	0.0261	0.0541***	0.0204
	(0.0107)	(0.0253)	(0.0134)	(0.0261)	(0.00731)	(0.0192)	(0.00808)	(0.0193)
b[Matchers] -b[Net-Givers]	p=0.134	p=0.139	p=0.291	p=0.296	p=0.898	p=0.346	p=0.713	p=0.424
b[Matchers | IPP] -b[Net-Givers | IPP]	.	.	p=0.085	p=0.079	.	.	p=0.002	p=0.005
Socio Demographics	X	✓	X	✓	X	✓***	X	✓***
Company Controls	X	✓	X	✓	X	✓	X	✓
Career Dummies	X	✓	X	✓	X	✓***	X	✓***
Dep. Dummies	X	✓	X	✓	X	✓***	X	✓**
Observations	846	831	836	831	863	857	863	857
R^2	0.007	0.046	0.025	0.054	0.006	0.236	0.018	0.244
Model	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS

Table 1.4 Regressions of Financial Appreciation on Cooperative Attitudes

Notes: Standard errors clustered on team-level in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; For wage increases, we use data from 2016/2017 for participants in 2017 and data from 2017/2018 for participants in 2019. We use the value of financial award payments received in 2017 in percent of the 2017-wage level. Asterisks for the control variables show the test result from an F-Test, testing the joint difference from zero.

ing or compensating differentials against the financial disincentives that may arise from cooperation when wage increases or merit-based awards are lower than for those who cooperate less.

In the following, we consider this mechanism and further plausible mechanisms that may be prima facie in line with our main results. In discussing alternative mechanisms, we do not necessarily assume that our three main results, (i) high cooperation levels, (ii) a negative nexus between cooperative attitudes and financial rewards, and (iii) a positive nexus between cooperative attitudes and non-financial outcomes, are connected. Obviously, only exogenous variation in some variables can provide a final answer on the sole driver of our results. However, some variables will never be varied exogenously in a meaningful way such as wage levels or wage increases. There is always a tradeoff between searching under the lamppost (and accepting that one studies very special setups that allow for exogenous variation) or using real-world environments that limit opportunities to exogenous variation. Nonetheless, we can provide heterogeneity analyses and robustness checks to shed light on the potential relevance of various mechanisms for our setting and for being in line with our main results.

1.6.1 High Levels of Cooperation

Our measures of cooperation are qualitatively comparable to the standard conditional contribution patterns documented in the behavioral economics literature; yet they appear higher (e.g., Fischbacher et al., 2001; Fischbacher and Gächter, 2010; Kocher et al., 2015).[22] To what extent cooperation rates in our setting reflect a general high level of cooperativeness of employees or rather a stronger role of a potential social desirability bias in our setup is a question that deserves further attention.

Within the framework of our study, we implemented an unexpected option to donate the experimental income at the end of the experiment in 2017. Participants could choose between receiving their income from the experiments in their personal bank account and donating it to one of five charities of their choice. At this point, participants did not know their income yet. We find a positive but insignificant relationship between donations and

[22] With respect to other non-student samples, Charness and Villeval (2009) observes that employees in the manufacturing industry contributed between 32% and 38% of their endowment to a three-person public good. Another example is Burks et al. (2016), who classify 24% of truck drivers in the same company as free-riders using a Prisoner's Dilemma Game. Algan et al. (2013, 2014) conducted public goods games with programmers at Sourceforge.net (an open source software platform) and users that contribute to Wikipedia, respectively. In both samples, subjects have already selected in a voluntary contribution platform; still, they are less cooperative, on average, than employees in our company (the 850 Sourceforge.net users unconditionally contribute 64% of their 10 tokens; the 1,194 Wikipedia users are less likely unconditional contributors and more likely free-riders than employees in our setting).

	(1) Donation	(2) Donation	(3) Donation	(4) Donation
Uncond. Contribution	0.0235			
	(0.0205)			
Belief About Others' Uncond. Contribution		-0.00976		
		(0.0213)		
Mean Cond. Contribution			0.0290	
			(0.0275)	
Net-Taker				0
				(.)
Matcher				0.156
				(0.160)
Net-Giver				0.259
				(0.161)
Constant	-0.283	-0.0318	-0.250	-0.259**
	(0.173)	(0.155)	(0.157)	(0.128)
N	438	438	438	438
Pseudo R2	0.002	0.0003	0.002	0.004
Model	Probit	Probit	Probit	Probit

Table 1.5 Regressions of Donations on Public Goods Game Measures

Notes: Robust standard errors in parentheses. Regressions are based on employees participating in the experiments in 2017; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

our public goods game variables (contributions and more cooperative types). This holds regardless of whether we use unconditional, conditional contributions or cooperative attitudes as regressors (see Table 1.5). Thus, donations seem to draw on a distinct concept than cooperative attitudes and cooperation. We consider this suggestive evidence that social desirability is not too much of an issue in our setup. Donation behavior (or dictator game giving more generally) is often thought of as being heavily affected by social desirability concerns. If cooperation in the public goods game was affected by social desirability concerns as well, we would observe a significantly positive correlation between the two sets of decisions.

In 2019, we implemented an additional public goods game after the main experiment in which the MPCR was set to either 0.3 or 1.2. Participants that are driven by social desirability concerns should be less likely to adjust their unconditional contribution to the reduction in the MPCR from 0.5 to 0.3, because they might want to signal cooperativeness. Responses to the increase of the MPCR to 1.2 should reflect mainly a sound understanding of the game's incentives. We elicited unconditional contributions, beliefs, and conditional contribution schedules for both alternative MPCRs, using the strategy method. We observe strong reactions to the two variations. Subjects significantly decrease unconditional contributions, beliefs, and conditional contributions when the MPCR decreases to 0.3 (means: 3.71, 2.91, 3.82, respectively, using Wilcoxon Signed-Ranks tests in comparison to the standard MPCR of 0.5; all p-values< 0.001). The reverse happens when the MPCR increases to 1.2 (8.82, 8.53, 8.37, respectively, using Wilcoxon Signed-Ranks tests in comparison to the standard MPCR of 0.5 all p-values< 0.001). We conclude that neither social desirability nor confusion are convincing explanations for the high levels of cooperation that we observe.

1.6.2 Negative Nexus Between Cooperative Attitudes and Financial Rewards

Following Bowles et al. (2001) and Barr and Serneels (2009), the correlation between cooperative attitudes and financial rewards could also be explained by an omitted variable bias with respect to skills that are specific to cooperative attitudes and related to performance differences. For example, Net-Givers could have a comparative advantage in networking or socializing and Net-Takers could be more strategically sophisticated. Table 1.6 shows OLS regressions of financial rewards on cooperative attitudes estimated for the two business models that exist in the company. As cooperation is more important in cloud-related jobs, we expect Net-Givers to perform better than Net-Takers in such jobs and thus receive higher wage increases or financial awards. However, Net-Takers receive significantly higher financial rewards than Net-Givers and Matchers (see columns (1) and (2) and column (5) and (6), respectively). This relationship does not exist in customer-related jobs (see columns (3) and (4) and columns (7) and (8), respectively). Thus, even if Net-Givers work on tasks with complementarities for which they should have the more appropriate cooperative attitude, Net-Takers get 2.1%-points higher annual wage increases and 2.5%-points higher award payments. The result indicates that there are no strong comparative skill and performance differences between attitudes; however, it might still be the case that Net-Takers have an absolute skill advantage. However, this would require Net-Takers, i.e. less pro-social types, to have, in general, higher levels of skill.

Another potential mechanism could be related to selection based on cooperative attitudes. Net-Takers could select into jobs with higher financial rewards, while Net-Givers could select into jobs with higher non-financial rewards (Falk and Heckman, 2009; Dohmen and Falk, 2011). Conversely, along the lines of Bowles (1998) and Levitt and List (2007), financial and non-financial rewards could also shape cooperative attitudes. Pay scheme specific norms could render selfish behavior in the individual performance scheme and pro-social behavior in the company performance scheme more appropriate and hence employees that comply with the norm get financially rewarded.

In line with both explanations, Appendix A.5 shows that employees in the individual per-

	(1) Wage Increase (in %)	(2) Wage Increase (in %)	(3) Wage Increase (in %)	(4) Wage Increase (in %)	(5) Fin. Award Payment (in % of Wage)	(6) Fin. Award Payment (in % of Wage)	(7) Fin. Award Payment (in % of Wage)	(8) Fin. Award Payment (in % of Wage)
Net-Takers	0	0	0	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Matchers	-0.0326**	-0.0230**	-0.00272	-0.00141	-0.00515	-0.0136**	0.00188	-0.00475
	(0.0143)	(0.0117)	(0.0138)	(0.0156)	(0.00764)	(0.00643)	(0.00899)	(0.00844)
Net-Givers	-0.0337**	-0.0213*	-0.00532	-0.00782	-0.0130*	-0.0245***	0.00164	-0.00453
	(0.0148)	(0.0123)	(0.0118)	(0.0117)	(0.00719)	(0.00701)	(0.00876)	(0.00920)
Constant	0.0672***	0.0934**	0.0303*	0.0505*	0.0655***	0.0680*	0.0340***	-0.0487*
	(0.0145)	(0.0457)	(0.0173)	(0.0278)	(0.0115)	(0.0357)	(0.0123)	(0.0263)
b[Matchers] -b[Net-Givers]	p=0.895	p=0.852	p=0.827	p=0.612	p=0.186	p=0.044	p=0.975	p=0.943
Included	Cloud	Cloud	Customer	Customer	Cloud	Cloud	Customer	Customer
Socio Demographics	X	✓	X	✓**	X	✓*	X	✓*
Company Controls	X	✓	X	✓	X	✓	X	✓***
Career Dummies	X	✓	X	✓	X	✓***	X	✓***
Dep. Dummies	X	✓	X	✓	X	✓***	X	✓***
Observations	410	406	236	233	422	417	248	245
R^2	0.022	0.109	0.005	0.129	0.008	0.294	0.042	0.294
Model	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS

Table 1.6 Regressions of Financial Appreciation on Cooperative Attitudes by Business Model

Notes: Standard errors clustered on team-level in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; For wage increases, we use data from 2016/2017 for participants of 2017 and data from 2017/2018 for participants in 2019. We use the value of financial award payments received in 2017 in percent of the wage level in that year in 2017. We exclude employees that neither work in the cloud nor in the customer business model. Here, we do not observe statistically relevant differences in financial awards with respect to cooperative attitudes. In terms of wage increases, we observe that Matchers receive slightly higher increases than Net-Takers which is marginally significant (p=0.071). Asterisks for the control variables show the test result from a F-Test, testing the joint difference from zero.

formance pay scheme are significantly more likely to be Net-Takers than employees in the company performance pay, which goes with conventional wisdom. Individual performance incentives do not foster pro-self behavior or they do not seem to attract more pro-social employees. Also, our survey analysis confirms that employees in the individual performance pay consider cooperation to be less important to fulfill their tasks successfully. If this observation is due to the idea that incentives shape preferences, we would expect that employees who already work for several years in the company and presumably in the same pay scheme exhibit pay scheme specific norms more strongly. Thus, we expect that employees get less cooperative the longer they work in the individual performance pay scheme. Table 1.7 shows results from an OLS regression assessing the effect of seniority on the relationship between cooperative attitudes and pay schemes. While the significant difference in mean conditional contributions between pay schemes remains, we find no significant interaction effect with seniority. This evidence suggests that there is a potentially stronger role for selection.

	Mean Cond. Contribution
Ind. Perf. Pay	-0.906**
	(0.407)
Seniority	0.012
	(0.0172)
Ind. Perf. Pay *	-0.003
Seniority	(0.0280)
Age	0.006
	(0.0131)
Stability	0.037
	(0.127)
Constant	4.704***
	(0.804)
Socio-Demographics	✓
Company Controls	✓
Career Dummies	✓
Dept. Dummies	✓
Observations	857
R^2	0.036
Model	OLS

Table 1.7 Regressions of Cond. Contributions on Pay Schemes and Seniority

Notes: Robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Asterisks for the control variables would show the test result from a F-Test, testing the joint difference from zero. For none of the variables the joint difference from zero can be rejected at conventional significance values. The variable *stability* capture the employee's feeling of team stability that incorporate a survey item on how long the employee works in his/her job function and hence pay scheme.

1.6.3 Positive Nexus Between Cooperative Attitudes and Non-Financial Rewards

We observe that the positive relationship between cooperative attitudes and sending awards appears widely insusceptible to context factors like pay schemes and business models. Based on simple regression similar to those used in Table 1.3, we observe that Net-Givers send significantly more recognition awards than their colleagues in the cloud (means per employee: 0.32 versus 0.12, $p = 0.085$) and the customer business model (means per employee: 0.65 versus 0.13, $p = 0.002$) as well as in the company performance pay scheme (0.41 per employee versus 0.13 per employee, $p = 0.004$). We also observe a similar pattern in the individual performance pay scheme that is however not statistically significant; admittedly, there is a relatively small sample size for this comparison (0.27 per employee versus 0.19 per employee, $p = 0.492$). The existence of the relationship across different company contexts suggests a more general link between recognition awards and cooperative attitudes, corroborating our external validity argument.

At the same time, we observe strong differences in reception rates between context factors. We find that reception rates are generally higher in the cloud model than in the customer-based model (0.31 per employee versus 0.21 per employee; $p = 0.046$) and in the company performance pay versus the individual performance pay (0.30 per employee versus 0.16 per employee; $p = 0.023$). This indicates that the recognition tool is used more frequently in areas in which teamwork and cooperation is required.

In our post-experimental survey, we elicit further variables that may relate to non-financial rewards or non-financial costs of cooperation. On the individual level, we capture work-related stress and overall work satisfaction. While our stress measure appears to be unrelated to conditional contributions (Spearman Correlation= -0.098, $p = 0.438$), we observe a strong positive correlation between cooperativeness of employees and work satisfaction (Spearman Correlation= 0.916, $p = 0.014$) that is robust to including personal and company controls. On the team-level, we measure perceived team cohesion and team stability. In Appendix A.9, we show that there exists no statistically relevant relationship between team stability and the share of Net-Givers in a team, but teams that perceive themselves as being more cohesive tend to consist of more Net-Givers.

1.7 Conclusion

This chapter provides novel evidence on how cooperative attitudes of employees are related to professional behavior and rewards within a large company. We observe high levels of cooperation among employees and evidence on the external validity of our experimental measure of cooperative attitudes for the company setting. In addition, we document a robust negative nexus between cooperative attitudes and financial appreciation, and a positive nexus between cooperative attitudes and non-financial rewards.

In line with a recent literature that emphasizes the intrinsic nature of cooperation (e.g., Hamilton et al., 2003; Bandiera et al., 2005, 2011, 2013; Ruff and Fehr, 2014) our analyses suggest that the company studied here positively affects levels of cooperation – despite financial disincentives for cooperators – through providing cooperative employees with non-financial compensations. We also document a potential role of selection based on cooperative attitudes in pay schemes similar to Burks et al. (2009).

Our findings have implications for the optimal design of incentives and management practices in companies that want to foster cooperation. A general implication is that companies should create a work context that allows non-monetary forms of rewards as values for cooperation to unfold. This might entail the opportunity for employees to voluntarily select into differently composed teams or work organizations, or the selection into organizational units with different cooperative cultures (Kosfeld and Von Siemens, 2011). At the same time, our findings stress the importance of management practices that operationalize the non-monetary returns of cooperation (like the recognition award systems used in our company).

We see our study as a first step and encourage other researchers to study cooperation in corporations as well. Obviously, we have no way to take firm conclusions regarding company-specific and more general results, given that our focus is on one company. It might well be that the specific interplay between incentives and culture at our company is different than in other companies. It might well be that the industry that our company is operating in has specific characteristics in terms of how cooperation is rewarded. Given the importance of cooperation in teamwork, it is astonishing that there is not more research empirically addressing the relationship between corporate culture, financial and non-financial rewards, and cooperation within the company. Although we believe that the gist of our results will hold more generally, given its systematic pattern, our results at the very least provide a proof of concept: The experimentally elicited measures on cooperation are systematically related to outcomes in the company. Our tests for external validity provide promising results.

We have searched for evidence outside the light of a lamppost, in contrast to some other studies that use more artificial designs in the wild to get more powerful inference. Both approaches seem useful. Next to understanding the causal mechanisms underlying our findings, a deeper understanding of the nature of the relationship between financial and non-financial incentives for cooperative behavior in organizations is required. Can financial and non-financial incentives work as substitutes on the individual employee level and, at the same time, work as complements when regarding the company's profits? How can the optimal mix of financial and non-financial incentives be characterized? More research is needed to empirically understand the optimal balance between cooperation-enhancing and competition-enhancing policies within organizations, probably dependent on cooperation culture and workforce composition.

Chapter 2

Cooperation, Free-Riding, and the Signaling Value of Incentives: An Experiment in a Company

2.1 Introduction

Complementarities in production render cooperation among employees important for companies (e.g., Dirks and Ferrin, 2001; Gratton, 2009). At the same time, they cause free-rider problems due to a misalignment of individual profits and collective efficiency (e.g., Gittell, 2000; Fehr, 2018). Using monetary incentives is a prevalent strategy of companies to cope with such conflict, but their effectiveness is still at debate.[1] Recent research points out that incentives can induce unintended side effects that eventually impede their original purpose (Gneezy et al., 2011; Bowles and Polanía-Reyes, 2012).

One effect that is of particular relevance for the context of cooperation is that incentives convey information about typical behavior of others (e.g., Sliwka, 2007; Van der Weele, 2012; Bénabou and Tirole, 2011). A manager who introduces incentives to cooperate may signal that employees would act selfishly otherwise. As a result, employees may expect less cooperative behavior from their colleagues and, in line with evidence on conditional cooperation (Fischbacher et al., 2001), cooperate less themselves.[2] Evidence from the lab-

[1] Examples include the introduction of manager guidelines that outline cooperative behavior as a requirement for promotion and salary increases, or the provision of peer-to-peer recognition tools in which employees can confer monetary awards to cooperative colleagues. See Gratton (2009) and `www.blog.bonus.ly/a-look-at-googles-peer-to-peer-bonus-system` for a description of how Google and British Petroleum implement these tools.

[2] The term "conditional cooperation" describes that people cooperate if they believe that others cooperate as well. There exists ample evidence about the prevalence of conditional cooperators in various samples (e.g.,

oratory suggests that employees potentially understand the signaling value of incentives (e.g., Galbiati et al., 2013), but field evidence is largely missing.

Studying the signaling value of incentives within companies is however difficult. Incentives and information about cooperative behavior held by managers are endogenous, and whether such information is available to managers might be unknown to employees. This study exploits a unique field environment that combines three very rare features that allow to overcome these issues. First, it allows for exogenous variation in information (about the cooperativeness of employees) held by managers when choosing incentives. Second, employees are well aware of the fact that cooperativeness measures exist. Third, employees also know whether such measures are (or are not) available to managers when setting incentives.

I collaborate with a large software company that relies heavily on cooperative behavior of their employees and seeks to provide incentives to encourage the latter. To study whether incentives work as signaling devices, I conduct an artefactual field experiment (Harrison and List, 2004) with managers and employees from the company. Employees find themselves in a social dilemma situation in which they have a dominant strategy to free-ride on the cooperative efforts of their colleagues. Managers benefit from high cooperation levels among employees and can counter free-riding by implementing a costly incentive that promotes cooperation. Prior to incentive choice, I exogenously vary whether managers are informed about prevailing cooperation levels among employees measured in a previous study (Deversi et al., 2020a). At the same time, I notify employees that their manager has been informed before setting incentives. By comparing beliefs and behavior of employees under informed versus uninformed incentive choices, I am able to isolate whether the information provided to managers transmits to employees and hence affects the company's cooperative culture.

I find that incentives have strong positive effects on cooperation. They increase cooperation rates by 24%, and beliefs about cooperative behavior of those working under incentives by 44%. I do not observe differential increases between the information treatments, neither in beliefs nor in actual behavior. This indicates that employees do not take into account the information conveyed by the managers' incentive choices. Unlike employees, managers react to the information that is made available to them. In the treatment group, they update their beliefs and, in line with maximizing their profits, choose incentives to increase cooperation less frequently.

It appears that the absence of a signaling effect is driven by the employees' misperception

Gächter, 2007; Kocher et al., 2008).

of the managers' decision-making. Employees do not expect managers to choose incentives based on their monetary benefits. Instead, they consider managers more likely to choose incentives when managers expect higher levels of cooperation. Hence, employees appear to interpret managers' choices to "reward" cooperation through incentive provision. An interpretation that relates to an important contextual factor for the effectiveness of incentives: the general relationship between management and employees.

My findings relate to a large influential literature in economics and management science dealing with the interaction of incentives and social preferences (for a review, see Bowles and Polanía-Reyes, 2012). According to this literature, incentives can crowd out prosocial behavior because they provide information about the person who sets the incentive, such as selfish intentions (e.g., Fehr and Rockenbach, 2003; Fehr and List, 2004) or his or her knowledge about the task (e.g., Bénabou and Tirole, 2003; Bremzen et al., 2015; Deserranno, 2019). Another channel to which this literature has alluded to is the signaling of principals' private information about social norms. In the experimental laboratory, Danilov and Sliwka (2017) investigate shirking behavior of agents that work on individual tasks under either fixed or variable pay contracts. They find an increase in agents' trustworthiness when the principal is informed about past effort provision and refrains from implementing a variable pay contract. Cardinaels and Yin (2015) show that the use of incentives to increase truthful behavior in a reporting task signals that other agents were likely to report dishonestly before. Both studies differ from my design by analyzing individual decisions rather than interactions of multiple agents.[3] Galbiati et al. (2013) use a two-agent minimum effort game and vary whether sanctions are endogenously set by an informed principal or exogenously set by the experimenter. They find that endogenous sanctions are more effective in enforcing high effort because they signal high effort provision in past rounds. My study makes a relevant contribution to this literature by providing a unique, naturally occurring test environment of signaling effects and their predicted adverse impact on cooperation. My results on the signaling hypothesis are particularly informative because they give rise to important contextual factors that render signaling and crowding out effects more or less likely to occur.

The remainder of this chapter is structured as follows. I first introduce the experimental design by describing my field setting and the experimental game. Then, I present the results in Section 2.3. In Section 2.4, I discuss potential explanations for the absence of the hypothesized signaling effect in my setting. Section 2.5 concludes this chapter.

[3] This implies that in both studies information about prevalent behaviors must affect agents' behavior via conformity preferences (Sliwka, 2007) or social esteem (Bénabou and Tirole, 2011) of agents, rather than trough reciprocity (Van der Weele, 2012) or effort complementarities (Friebel and Schnedler, 2011) as in my setting.

2.2 Experimental Design

2.2.1 Field Setting

This study is conducted in partnership with a large software company. In most tasks within the company - reaching from software development, consulting, sales to service activities (e.g., human resource management) - cooperation is essential to maximize joint production output of work teams.[4]

The management of the company conducted a study to measure the prevailing levels of cooperation and to subsequently establish new policies that enhance cooperation. This study is described by Deversi et al. (2020a). It entailed a one-shot, three-person public goods experiment in which a total of 369 employees participated.[5] The data revealed high levels of cooperation (on average 79% of the endowment) and high expectations about others' cooperation behavior (on average 66% of the endowment) that were however significantly lower than actual cooperation rates. Further, about 82% of company employees were conditional cooperators which emphasizes the relevance of beliefs about others' behavior for cooperation in the company. Both results together indicate a significant room for signaling effects to adversely affect the cooperative culture of the company. If the management was to implement incentives without informing employees about the results of Deversi et al. (2020a), employees might infer that measured cooperation levels were low. The experiment of the current study takes place after the previous study, but before managers and employees have been informed about the findings.

2.2.2 Experimental Game

In the experiment, three randomly grouped employees ($n = 3$) play a public goods game. Each employee receives an initial endowment of 10 Tokens (worth € 10) to be allocated between a private account and a common account. The amount contributed to the common account is an integer that satisfies $0 \leq c_i \leq 10$. The sum of contributions to the common account is multiplied by 1.5 and then divided equally among the three group members. Therefore, each individual group member receives a share of $\gamma = 0.5$ of the total sum of contributions.

[4]For more detailed information on the company see Deversi et al. (2020a).

[5]The authors use a linear public goods game - also known as voluntary contribution mechanism. The incentives of the game capture a tension between individual payoff mazimization and collective efficiency maximization. In the game each player has a dominat strategy to free-ride on others' contributions to a public good, deviations from this strategy are usually interpreted as cooperative behavior or as a social preference more generally (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002).

In addition, I match one manager to each group of employees. They earn a fixed amount of 15 Tokens (worth € 15) and a share of $\gamma = 0.5$ from the sum of contributions. They cannot contribute. Before employees act, managers decide whether to implement a monetary incentive to make employees cooperate (termed *Additional Payment* in the instructions). If the incentive is chosen, the employee with the highest contribution in the common account receives an additional payment of three tokens.[6] Setting the incentive to cooperate ($a = 1$) costs 5 Tokens.

The payoff functions of employee i and manager m can be described as follows.

$$\pi_i = 10 - c_i + \gamma \sum_{j=1}^{n} c_j + a \times (\mathbf{1}\{c_i > \boldsymbol{c}_{-i}\}) \times 3 \tag{2.1}$$

$$\pi_m = 15 + \gamma \sum_{j=1}^{n} c_j - a \times 5. \tag{2.2}$$

where $\boldsymbol{c}_{-i}$ is the contribution vector of the other two group members.

If the incentive is not implemented ($a = 0$), the standard social dilemma equilibrium arises as $1/n < \gamma < 1$, i.e., it is welfare-efficient if each member contributed his or her whole endowment but individually optimal to contribute $c_i = 0$. If the incentive is implemented ($a = 1$), the dominant strategy depends on the expectation about others' contributions. For expected average contributions $E(\bar{c}_{-i}) \in [0, 5)$, it is payoff-maximizing to contribute $k = \min\{n \in \mathcal{N} | n > E(c_{-i})\}$, i.e., the minimal integer higher than $E(c_{-i})$. For $E(\bar{c}_{-i}) = 5$, the employee is indifferent between free-riding or contributing 5. For $E(\bar{c}_{-i}) > 5$, the social dilemma equilibrium emerges again. Overall, the incentive increases the expected payoff from contributing into the common account without affecting the action space of players.

From the managers' perspective, implementing the incentive can only be payoff maximizing if the expected sum of contributions without the incentive is lower than 20 Tokens (i.e., 6.67 Tokens per employee). In order for the cost of the incentive to pay off, each group member must increase contributions in response to the incentive by at least 3.33 Tokens.[7]

[6] The tie-breaking rule is specified such that the three tokens are evenly distributed among the participants that contributed the highest amount. I focus on this particular incentive because it is a policy that the management discussed to implement after conducting the analyses in Deversi et al. (2020a). The idea was to introduce a tournament incentive that rewards the employee with the highest number of received peer-to-peer recognition awards that can be sent in the companies intranet. Similar relative rewards for cooperation have been analyzed by Irlenbusch and Ruchala (2008).

[7] To see this, I compare the manager's payoffs $\pi_m(a = 1) - \pi_m(a = 0) = \gamma E[\sum c_j(a = 1) - \sum c_j(a =$

For both choices of the manager (i.e., using the strategy method), I elicit three decisions from the employees. First, I elicit their contribution in the common account (*unconditional contribution*). Second, I ask for their contributions if the other group members contributed on average 0/1/2/.../10 (*conditional contributions*). For one randomly selected subject in the group the conditional contributions are payoff-relevant, whereas for the two remaining subjects the unconditional contribution is. This ensures that both unconditional and conditional contribution decisions are incentive-compatible. Third, I elicit their belief about the average unconditional contribution of the other two players (*belief*). Following Gächter and Renner (2010), employees receive €5 if they hit the correct average, and €0 otherwise.

Finally, I ask two further questions that capture employees' beliefs about managers' incentive choice and their beliefs about managers' expectation about contribution behavior of employees. Both questions are incentivized by providing €1.5 for a correct response. A full list of elicited variables, including additional survey variables, can be found in Appendix B.1.

2.2.3 Treatments and Hypotheses

The critical feature of my experiment is the information structure. Generally, there exists uncertainty about employees' behavior in the game. I provide information on average unconditional contributions measured by Deversi et al. (2020a) to managers in Info, but not in No Info. Prior to incentive choice, they receive the following information.

> "*Tip for you as a manager: 369 employees have already made their decision to allocate the 10 tokens between the private account and the common account. There was no additional payment for these decisions in place. On average, 2.10 Tokens were paid into the private account and 7.90 Tokens into the common account.*"

On the employee side, the instructions in Info entailed the following statement.[8]

> "*What does the manager know before making a decision? The manager received information about the average contribution decision of 369 other employees.*

$0)] - 5 \geq 0$. Re-formulation yields $E[\sum c_j(a=1)] - E[\sum c_j(a=0)] \geq 10$, hence, $\frac{10}{3}$ per group member. In addition, as $\max E[\sum c_j(a=1)] = 30$, this yields an upper bound for the expected sum of contributions without the incentive, i.e., $E[\sum c_j(a=0)] = 20$.

[8] For employees in Info, the treatment information was referred to three times: once in the main instruction text, once on a summary screen with the most important aspects in bullet points, and another time in the comprehension tasks section where I asked a question on whether the manager has been informed.

These employees have already decided on the allocation of the 10 tokens between the private account and the common account. There was no additional payment for these decisions in place."

		Within Subject	
		No Incentive	Incentive
Between Subject	No Info	205 employees & 23 managers	
	Info	196 employees & 24 managers	

Table 2.1 Treatment Overview

Table 1 summarizes the design. It enables me to observe beliefs and cooperation of employees under different information sets of the managers while holding incentive choices constant. To derive testable predictions, I assume that both players update their beliefs in a Bayesian fashion and that managers are individual payoff-maximizers whereas employees are either individual payoff-maximizers as well or conditionally cooperative. Conditional cooperators contribute to the common account if they believe that others contribute as well.

Under these assumptions, managers should update their prior beliefs according to the average contribution rate provided to them in the information condition. They should respond to this belief update by choosing the incentive less frequently as measured contribution rates are higher than the critical threshold for choosing the incentives (7.90 Tokens $>$ 6.67 Tokens).

Hypothesis 1: *On average, managers update their beliefs according to the information provided and select the costly incentive less frequently in Info than in No Info.*[9]

The incentive should steer selfish employees away from free-riding. This should render employees' beliefs in others' cooperativeness more optimistic and further enhance contributions of conditional cooperators.

Hypothesis 2: *On average, employees' beliefs about others' contributions and actual contributions are higher in Incentive/No Info than in No Incentive/No Info.*

Consider employees' responses in No Incentive/Info versus No Incentive/No Info. Here, a manager decided not to intervene and the public goods game is played without the incentive to cooperate. In the Info treatment such choice reflects that contribution levels observed by the manager have been sufficiently high, as otherwise it would have

[9] There also is a equilibrium effect at work such that managers anticipate signaling effects from their incentive choices and hence choose the incentive even less frequently.

been worth to incur the cost to implement the incentive. Conversely, Incentive/No Info versus Incentive/Info should reflect the information that contribution levels observed by the manager have been sufficiently low, such that it was worth it to incur the cost to implement the incentive.[10] These belief updates should affect contribution behavior to the extent that employees are conditionally cooperative.

Hypothesis 3: *On average, employees' beliefs are more optimistic in No Incentive/Info compared to No Incentive/No Info and more pessimistic in Incentive/Info compared to Incentive/No Info.*

Hypothesis 4: *On average, employees' contributions are higher in No Incentive/Info compared to No Incentive/No Info and lower in Incentive/Info compared to Incentive/No Info if they are conditionally cooperative.*

2.2.4 Procedures

This study is part of a larger research agenda taking place in the company such that the experimental procedures that I used are identical to those described in Deversi et al. (2020a). Participants were randomly selected from a large population of employees eligible to participate in experiments that were taking place at the same time.

I conducted the experiment in spring 2019 using the software Qualtrics.[11] Potential participants were invited via e-mail and participated through a personalized link. Participation took place in a two-week time period. Payout calculations and matching of managers and employees were administered ex post. While there was no feedback during the experiment, participants received payoff information afterwards via a website created solely for this purpose. I asked participants to perform all experimental tasks individually and groups were randomly allocated to avoid coalition formation. A double-blind data procedure ensured the anonymity of all managers and employees. Approval of the ethics committee at the University of Munich has been granted in January 2019 and my analyses have been pre-registered at the AEA RCT registry (AEARCTR-0003931).

[10] The manager's actual decision threshold might be lower, depending on managers' beliefs and reciprocity preferences of employees (see Van der Weele, 2012), the upward containment is however unaffected by these other aspects. Hence, I expect employees to infer the positioning of the observed contribution levels relative to the upper threshold from managers decisions which implies that the empirical distribution of beliefs should shift.

[11] The instructions can be found in Appendix B.2.

	Managers		Employees	
	Mean	*Std. Dev.*	*Mean*	*Std. Dev.*
Female	0.40	0.50	0.33	0.47
Age	43.96	10.05	36.15	8.35
Seniority	11.73	6.97	5.08	3.89
Education				
Highschool	0.06	0.25	0.10	0.30
Bachelor	0.04	0.21	0.14	0.35
Master	0.63	0.49	0.60	0.49
Ph.D.	0.21	0.41	0.12	0.32
Other	0.06	0.25	0.04	0.19
Performance Pay				
Company			0.70	0.46
Individual			0.30	0.46
Observations	47		401	

Table 2.2 Participants' Characteristics

2.2.5 Sample Characteristics

I invited 1,500 managers and employees to participate in the experiment. A total of 48 managers and 401 employees participated which corresponds to a participation rate of 30%. Table 2.2 shows the sample characteristics.[12] Participating managers and employees are highly educated (only less than 14% have no post-secondary education). There are 19 female managers (40%) and 132 female employees (33%). Managers are on average 44 years old and work in the company for almost 12 years. Employees are on average 36 years old and work in the company for around 5 years. Furthermore, 70% of employees work under a company performance pay scheme in which bonuses depend on the company's asset market performance. The other 30% work under an individual performance scheme in which they receive bonuses based on individual target achievement. Many managers, especially those high in the hierarchy, have special contracts that can not be assigned to either of these schemes.

2.3 Results

2.3.1 Managers

A general prerequisite for a signaling effect is that managers react to the information treatment. Figure 2.1 shows the cumulative distribution function of the deviation between the managers' expectation and the average contribution level provided in Info. It becomes clear that managers hold heterogeneous beliefs in No Info that differ substan-

[12] The balance table is provided in Appendix B.3.

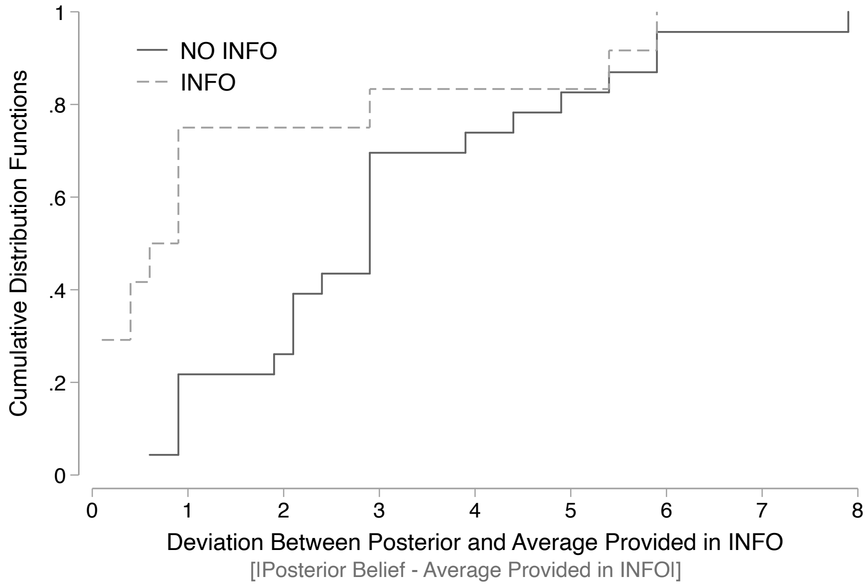

Figure 2.1 Treatment Effects on Managers' Posterior Beliefs

Notes: The graph shows the cumulative distribution functions of the absolute difference between managers' posterior beliefs about employees' contributions without the incentive in place and the measured contribution rate in Deversi et al. (2020a) by treatment.

tially from the provided average, and that managers in the information condition adjust their priors accordingly. Almost 80% of managers in Info deviate not more than one Token from the provided average value, whereas 20% hold such beliefs in No Info. Hence, a Mann-Whitney U Test (MWU) rejects that beliefs in both conditions are from the same underlying distribution ($p = 0.001$). This belief update should induce less selection of the incentive for profit-maximizing managers. And indeed, I observe that managers select the costly incentive less frequently in Info than in No Info (71% versus 91%). However, as I observe only 47 managers' decisions, this difference is only marginally statistically significant (MWU, $p = 0.078$; or one-sided Fisher Exact test, $p = 0.078$).

Result 1: *Managers' beliefs are significantly closer to the measured contribution rates in Info than in No Info, and managers select the incentive less often in Info than in No Info.*

2.3.2 Employees

As described in my hypotheses, beliefs about others' contributions are a crucial indicator for the mechanisms driving potential effects in the incentive and information conditions. Figure 2.2 presents the respective treatment comparisons. Beliefs about others' unconditional contributions are higher when the manager selected the incentive as compared to when it was not selected (7.5 Tokens versus 5.2 Tokens; Wilcoxon Signed-rank Tests

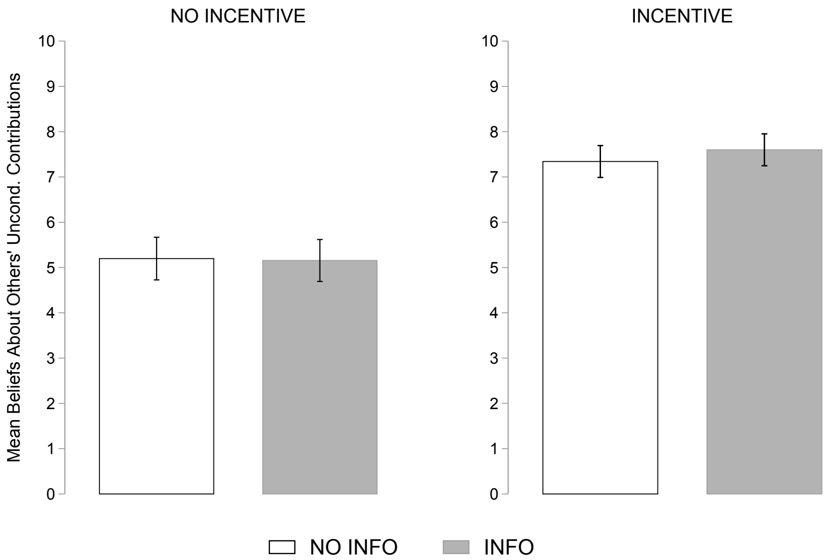

Figure 2.2 Treatment Effects on Employees' Beliefs

Notes: Bars show the average belief of employees about the unconditional contribution decision of the other group members. The 95% confidence intervals are based on a standard normal distribution.

(WSR), p< 0.001). This difference is also statistically significant when tested in both treatments separately (WSR, both p< 0.001). Yet, the information treatment has no impact on beliefs, neither under No Incentive (MWU, p= 0.906) nor under Incentive (MWU, p= 0.236). The individual within-subject difference in beliefs between the two incentive states is also not statistically significant from each other between Info and No Info (MWU, p= 0.314). This indicates that employees' beliefs were unresponsive to the information treatment.[13] If anything, we observe a small tendency in the opposite direction of the predicted effect.

To show a more complete representation of the belief data, Figure 2.3 plots the cumulative distribution functions of the individual belief differences between Incentive and No Incentive. If incentive choices work as signaling devices, the difference in beliefs should be lower in Info compared to No Info. However, I do not find an indication for this effect. Both distributions appear very similar to each other and do not clearly diverge (Kolmogorov-Smirnov Test, $p = 0.402$).

The estimation results in column (1) of Table 1 confirm the non-parametric analyses.

[13]This null result seems not to be driven by low statistical power. In my *ex ante* power analysis, I calculated a required sample size of 368. Considering the final sample size of 402, my experiments appear slightly overpowered and still show a null result. In the *ex post* power calculation, given my sample size and the measured standard deviations in the belief difference between the incentive states, I would be able to detect an effect size of 30% of a standard deviation which is smaller than detected effect sizes in, for example, Galbiati et al. (2013) or Cardinaels and Yin (2015).

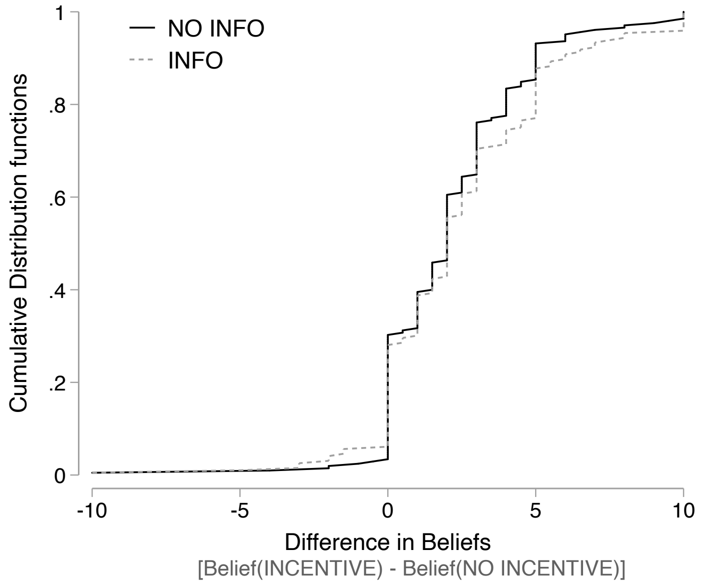

Figure 2.3 Treatment Effects on Employees' Beliefs About Incentive Effects

Notes: The graph shows the cumulative distribution functions of the difference between employees' beliefs about others' contributions with the incentive in place and versus without the incentive in place by treatment.

Here, I regress beliefs on treatment dummies. The OLS regression pools all decisions in the strategy method and uses clusters on the subject level. While the incentive significantly increases beliefs by 44% (2.1 Tokens) on average, the interaction of the information treatment and the incentive choice as well as the information dummy alone have only small positive and insignificant effects.

The null result of signaling effects on beliefs renders potential effects on behavior in the public goods game unlikely. Still, as beliefs were elicited after public good contributions, it might be the case that order effects biased belief updating but not potential effects on behavior. As presented in column (2) of Table 1, I observe however comparable effects on unconditional contributions. The incentive decision induces an increase in unconditional contributions by 23% (1.5 Tokens), but there is no statistically significant effect of the information treatment or the treatment interaction. Furthermore, the estimated models in columns (4) to (6) show that the null effect of the treatment interaction is robust to controlling for a wide range of employee characteristics including *gender*, *age*, *seniority*, *incentive scheme*, *career level*, and *job function.*

Result 2: *Employees' beliefs about others' contributions and actual contributions are significantly higher in* Incentive *than in* No Incentive.

Result 3: *Employees' beliefs are not statistically different between* No Incentive/Info *and*

	(1) Belief	(2) Uncond. Contribution	(4) Belief	(5) Uncond. Contribution
I(Incentive)	2.144***	1.346***	2.172***	1.376***
	(0.174)	(0.216)	(0.179)	(0.219)
I(Info)	-0.0419	-0.226	-0.0442	-0.168
	(0.335)	(0.363)	(0.339)	(0.364)
I(Incentive × Info)	0.300	0.391	0.294	0.329
	(0.274)	(0.319)	(0.283)	(0.326)
Constant	5.198***	6.744***	4.912***	6.552***
	(0.239)	(0.248)	(0.494)	(0.528)
Controls	No	No	Yes	Yes
Observations	802	802	784	784
R^2	0.131	0.055	0.156	0.092
Model	OLS	OLS	OLS	OLS

Table 2.3 Regression Estimations of Treatment Effects

Notes: For each employee and dependent variable two entries are observed: one entry under the incentive and one without the incentive. The control variables include *gender*, *seniority*, *incentive scheme*, *career level*, and *job function*. 18 employees are not included in the regressions using the additional controls as some of these have not been available for those participants. Standard errors are clustered on the subject level and are shown in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

No Incentive/No Info. They are also not statistically different between Incentive/Info and Incentive/No Info.

Result 4: *Employees' contributions are not statistically different between No Incentive/Info and No Incentive/No Info. They are also not statistically different between Incentive/Info and Incentive/No Info.*

2.3.3 Treatment Heterogeneity

Following Danilov and Sliwka (2017), one may expect that employees that work at the company for only a short period of time should update their beliefs more strongly because they have a less precise prior. In columns (1) and (2) of Table 2.4, I show OLS regressions for employees whose seniority is above and below the median seniority level, respectively. For less senior employees, the interaction effect of the incentive choice and the information treatment is positive and marginally significant. I.e., these employees exhibit a small tendency to infer relatively high cooperation rates from managers setting the incentive. For more senior employees, the interaction is very close to zero and insignificant.

With respect to cooperation behavior, one may expect hat employees with strong reci-

	(1) Belief	(2) Belief	(3) Uncond. Contribution	(4) Uncond. Contribution
I(Incentive)	2.329***	2.000***	1.369***	1.393***
	(0.235)	(0.279)	(0.278)	(0.331)
I(Info)	-0.778	0.453	-0.176	-0.303
	(0.498)	(0.472)	(0.433)	(0.667)
I(Incentive) $\times$ I(Info)	0.712*	-0.014	0.692*	-0.468
	(0.382)	(0.419)	(0.412)	(0.504)
Constant	6.187***	6.305***	6.523***	6.497***
	(1.245)	(1.177)	(0.573)	(1.084)
Subgroup	Low Sen.	High Sen.	Cmp. Pay	Ind. Pay
Controls	Yes	Yes	Yes	Yes
Observations	384	400	552	232
R^2	0.206	0.169	0.120	0.111
Model	OLS	OLS	OLS	OLS

Table 2.4 Treatment Effects on Beliefs and Contributions by Subgroups

Notes: For each employee and dependent variable two entries are observed: one entry under the incentive and one without the incentive. The control variables include *gender*, *seniority*, *incentive scheme*, *career level*, and *job function*. 18 employees are not included in the regressions using the additional controls as some of these have not been available for those participants. Standard errors are clustered on the subject level and are shown in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

procity preferences react more strongly to a belief update. Using data from the previous study, I observe that employees working under individual performance pay are less likely to be conditional cooperators than employees under company performance pay (MWU, $p = 0.028$). As shown in columns (3) and (4) of Table 2.4, I observe that for employees in the individual performance pay scheme the coefficient of the treatment interaction is negative whereas for employees in the company performance pay scheme the coefficient is positive and also marginally significant.

As I will show in the next section, the observation that signaling effects have a small tendency to work in the opposite direction of my prediction is related with the employees' perception about how managers make incentive choices.

Result 5: *Less senior employees and employees that work under the company performance pay scheme exhibit a small tendency to infer relatively high cooperation rates from managers setting the incentive.*

2.4 Discussion

Why is there no signaling effect of incentive choices? To begin with, a basic requirement for a causal treatment effect is that participants paid attention to treatment specific information and understood the incentive structure of the game. In this regard, it is affirmative that participants in Info took longer to complete the experiment than employees in No Info (MWU, $p = 0.005$). Next to reading the additional instructions, this might also entail some time in which employees were thinking about the implications of the managers being informed when setting incentives. Comprehension questions at the beginning of the experiment and a telephone hotline through which participants could ask questions during the experiment aimed at preventing misunderstandings.

Several potential explanations for the null evolve from the complexity of the reasoning process required from employees to infer signals from managers' choices. It should first be noted that the strategic sophistication of participants in my sample is arguably high. Compared to standard student subject pools, employees have a high average education level including many employees with a PhD. As math skills are often found to be positively related to strategic sophistication (e.g., Czermak et al., 2016), I asked participants how well they feel described by the statement "I am good at maths" in the post-experimental survey. I find that the median response on a scale from 0 ("does not describe me at all") to 10 ("describes me perfectly") is relatively high at 7. I also show in Appendix B.4 that my main regression results are robust to accounting for this variable. Moreover, it has been argued that using the strategy method to represent managers' choices "[...] may signal to agents that the experimenter wants them to infer information from contract choices" (Cardinaels and Yin, 2015, p. 1012); such a reflection effect essentially limits the strategic sophistication required from participants and makes my null result even stronger.

Even under the premise of high strategic sophistication, it might still be the case that stakes involved for employees were too low, i.e., participants did not spent the cognitive efforts required to process the conveyed information. However, Deversi et al. (2020a) find indications that employees from the company cared about similar-sized stakes in a public goods game. In their experiment, a substantial share of participants reacted to variations in the marginal per capita return of contributions in the common account. Also, in a surprise donation option at the end of their experiment most participants decided to keep the final payoff for themselves rather than donating it to a charity.

I now turn to considering employees preferences and their beliefs about managers' decision-making as a potential source of the null. First, it could be that employees are not conditionally cooperative such that they do not value the potential signals about others' behavior. Following Fischbacher and Gächter (2010), I estimate an individual reciprocity parameter for each employee using the conditional contributions in No Incentive/No Info.[14] I find that the average value of the parameter is 0.7 and that there exists a substantial fraction of perfectly conditional cooperators among employees (48%) who should care a lot about information about others' contributions.

Second, employees might expect managers to be indifferent between selecting the incentive or not such that managers' choices are random and do not signal. Contrary to this concern, I observe that on average employees expect managers to select the incentive with a likelihood of around 63% which is significantly different from 50% (WSR, $p < 0.001$) and almost identical between treatments (63.3% in Info and 63.4% in No Info; MWU, $p = 0.976$).

Third, it would be detrimental to the signaling effect if employees misinterpret the managers' purpose of providing the incentive. There is some evidence in line with this argument. Some employees (21%) expect their managers to expect zero or even negative incentive effects on contributions. Hence, in Table 2.5, I re-estimate the main OLS regressions from Table 2.3 excluding these employees. Interestingly, I observe that the positive interaction effects for beliefs and unconditional contributions increase compared to the full sample estimates. The signaling effect on beliefs is even statistically significant at the 5% level. Employees infer high contribution levels from informed managers that select the incentive which leads to a crowding-in effect on contributions that is marginally significant.

What reasoning do employees expect from managers that can explain these observations? In Figure 2.4a, I correlate the expected likelihood of the managers setting the incentive with employees' beliefs about the manager's expectation of the unconditional contribution levels.[15] If employees perceive the managers as individual profit-maximizers who tradeoff the expected incentive effect against its costs, one would observe a positive relationship between both variables. However, I observe that employees perceive them as independent (slope parameter in No Info of -0.01, t-Test, $p = 0.993$). The relationship turns slightly positive in Info but remains insignificant (interaction effect of 0.91, t-Test, $p = 0.450$). Employees appear to not take into account that setting the costly incentive fulfills a selfish purpose. In Van der Weele (2012) or Bénabou and Tirole (2011), for ex-

[14] If the parameter is 1, there is a linear relationship between an employee's contributions and the average contributions of the other two employees in the contribution schedule. The parameter is 0 if the employee's and the others' contributions are independent from each other.

[15] There are no significant differences in these second-order beliefs between Info and No Info (MWU, $p = 0.400$) corroborating the null result further.

	(1) Belief	(2) Uncond. Contribution
I(Incentive)	2.642***	1.580***
	(0.180)	(0.238)
I(Info)	-0.291	-0.319
	(0.340)	(0.403)
I(Incentive) × I(Info)	0.563**	0.668*
	(0.280)	(0.355)
Constant	4.559***	6.379***
	(0.251)	(0.277)
Excluded	Misperceivers	Misperceivers
Controls	No	No
Observations	640	640
R^2	0.220	0.085
Model	OLS	OLS

Table 2.5 Second Order Beliefs and Treatment Effects on Beliefs and Cooperation

Notes: For each employee and dependent variable two entries are observed: one entry under the incentive and one without the incentive. 81 employees are not included in the regressions as they do not expect that their managers expect higher cooperation from setting the incentive (i.e., they misperceive the purpose of the incentive) . Standard errors are clustered on the subject level and are shown in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ample, it is a necessary requirement for a signaling effect that the agents presume their managers to choose an incentive if it helps them to increase personal profits. Otherwise, the employees can not infer from incentive provision that contribution levels were low. The lack of evidence confirming this presumption can hence explain the overall null result. However, it does not explain the observed crowding-in tendencies in the data.

To analyze this, we need to understand how the presumption of employees about managers' choices actually looks like. In Figure 2.4b, I correlate employees' beliefs about the likelihood of setting the incentive with their beliefs about the managers' expectation of the unconditional contribution level without the incentive in place. While one would expect a downward sloping relationship in line with payoff maximization, I find the opposite. A standard deviation increase in the belief about the managers' expectation increases the belief about the likelihood of incentive selection by 2.5%-points (t-Test of regression coefficient, $p < 0.001$). It appears that employees expect that managers reciprocally provide rewards for high expected levels of cooperation. Employees do not see their managers as selfish profit maximizers. This could be related to past experiences with managers or a general prosocial relationship between management and employees in the company. As employees think that managers provide incentives based on high ex-

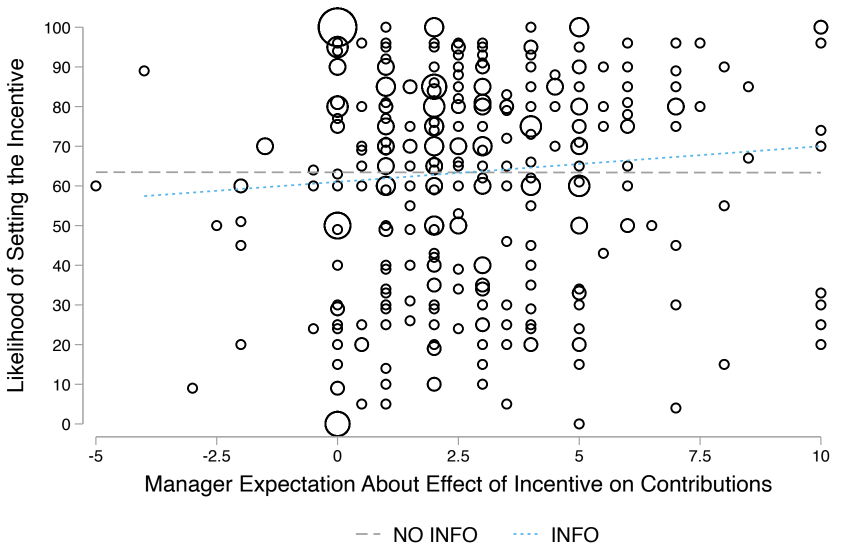

(a) Managers' Incentive Choice and Expected Incentive Effect

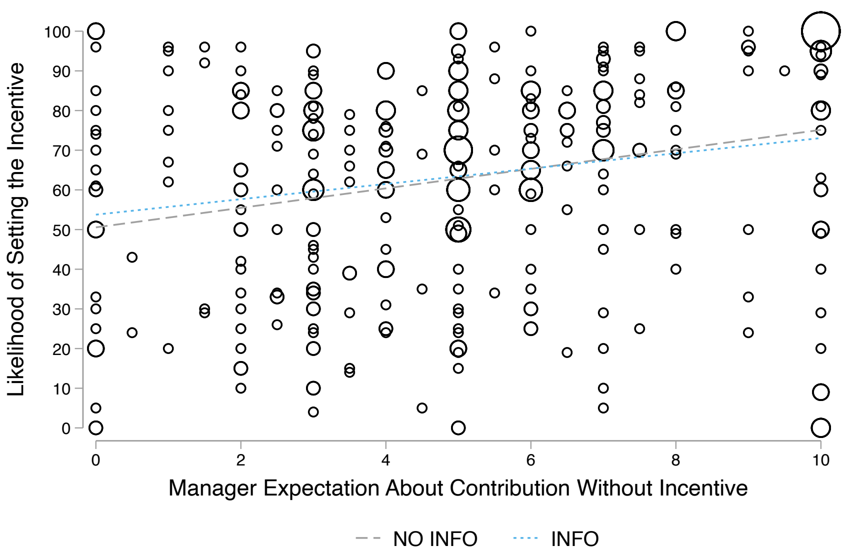

(b) Managers' Incentive Choice and Expected Contributions

Figure 2.4 Employees' Beliefs About Managers' Decision-Making pectations about cooperation rates, incentive provision signals high contribution levels and can explain the belief update observed in Table 2.5.

2.5 Conclusion

The literature suggests that incentives designed to promote cooperation in organizations may signal that selfish behavior is prevalent. As a consequence, they only have limited or even counterproductive effects. Contrary to this hypothesis, I find that setting an incentive to cooperate significantly increases cooperation among employees of a large software company. This increase is not affected by signals about others' behavior.

Further analyses suggest that the absence of a signaling effect in my setting is related to employees' perception of their managers' decision-making. They believe that managers do not exploit their private information about others' behavior in an opportunistic manner, but provide incentives if they expect high levels of cooperation. This might explain why I observe a small tendency in the data that employees infer high cooperation levels from incentives set by informed managers.

To the best of my knowledge, my study is the first to analyze whether contract choices signal social norms in a relevant field environment. According to Levitt and List (2007), it is often not possible to generalize findings from the experimental laboratory to the field because contexts differ. Actors in the field bring internalized social norms or past experiences and strategies into the game and herewith change outcomes. In my partner company and probably other organizations alike, reputation appears to be an important context factor of the functioning of incentives. A more nuanced understanding of this and other contextual factors, for example, the transparency about superior information on the side of the principal or the legitimacy of principals' decision making (Schnedler and Vadovic, 2011), is required. Another question for future research that arises from my setting is whether companies can prevent signaling effects of incentives by actively investing in the general relationship between managers and employees. This might include establishing pro-social intentions in managers such that their decision making "serves the employees", or to create a perception among employees that the management pursues benevolent management strategies.

Finally, it must be noted that in most field experiments there exists a tradeoff between using more artificial designs to discover causal effect mechanisms underlying the data and more natural designs that allow for bigger picture analyses (Deversi et al., 2020a). This chapter focused on teasing out the signaling of others' behavior via incentive choices. Companies that design incentives to promote cooperation should also take other forms of incentive effects, like framing effects or the signaling of other information hold by the management (Bowles and Polanía-Reyes, 2012), into account.

Chapter 3

Spin Doctors: An Experiment on Vague Disclosure*

3.1 Introduction

In many settings, informed parties not only decide whether to disclose verifiable private information, but also enjoy substantial flexibility in how information is disclosed. One way to exploit flexibility in disclosure is by means of vague messages. Vague messages are designed to inflate a receiver's perception of the sender's type by clearly separating from worse but not from better types. They are not outright lies, which may invite litigation, but merely put a positive spin on unfavorable news. Consider the following examples.

A college that ranks 10th in the latest US news ranking is likely to call itself a top 10 college rather than referring to itself as the 10th ranked college. A wine whose sole designation of origin is France is unlikely to come from the Bordeaux region, renowned for its superior wine. A wine whose sole designation of origin is Bordeaux is unlikely to come from Pomerol, an especially beloved subregion of Bordeaux. Researchers often refer to "significance at the 5 percent level" when a p-value is just below 0.05, while stating the exact p-value for a highly significant result. During legal proceedings, a defendant may try to convince a jury of her innocence by answering only those questions that are likely to exonerate her.

Sophisticated receivers understand and can correct for senders' strategic use of vagueness. But if these deceptive practices are deployed on naive receivers, then they result in systematic misperceptions. We model voluntary disclosure to receivers of heterogeneous strategic sophistication under both flexible language, which facilitates vague messages,

*This chapter is based on joint work with Alessandro Ispano and Peter Schwardmann.

and precise language. We then test the model's assumptions and predictions in the experimental laboratory. In doing so, we seek to answer three main questions. How do senders optimally design messages to exploit receivers' naivete? Are (some) receivers systematically fooled by vague disclosure? And can restricting senders' flexibility in disclosure improve information transmission?

The Model. In order to derive behavioral predictions for the experiment we adopt a model due to Milgrom and Roberts (1986), Eyster and Rabin (2005) and Hagenbach and Koessler (2017). To allow us to speak to the model's policy implications, we derive additional welfare results.

Consider a voluntary disclosure game in which a privately informed sender decides whether and how to disclose verifiable information about her type to a receiver. The sender's payoff is increasing in the receiver's belief about the sender's type, while the receiver's payoff is increasing in the accuracy of her belief. We distinguish between two language regimes: in the precise language regime, if the sender discloses, then the message has to reflect her exact type; in the flexible language regime a sender may send vague messages, i.e. a message that is any interval that contains the sender's true type.

If all agents are rational, in both the precise and the flexible language regimes the equilibrium features full information revelation (Grossman and Hart, 1980; Grossman, 1981; Milgrom, 1981).[1] However, the arguments for full information revelation and the irrelevance of language crucially depend on a high degree of strategic sophistication on behalf of the receiver. In reality, many receivers may be naive and struggle to be maximally skeptical in the face of nondisclosure or vague messages. Building on Milgrom and Roberts (1986), Eyster and Rabin (2005) and Hagenbach and Koessler (2017), our model therefore features both sophisticated and naive receivers. When a naive receiver encounters nondisclosure, she estimates that the sender is the average type. When she encounters a vague message, she estimates that the sender's type is the average of the sent interval.

The presence of naive receivers drives both nondisclosure (under precise language) and the exploitative deployment of vague messaging (under flexible language). Vague messages take the following simple form. Senders send an interval that spans their actual type and the upper bound of the message space.

Moving from the flexible language regime to the precise language regime then implies

[1]In the precise language regime, the highest type discloses because the disclosed information definitely exceeds receiver expectations. Because nondisclosure now cannot stem from the highest type, the second highest type is compelled to disclose. An iteration of this reasoning yields full disclosure. In the flexible language regime, the receiver's belief that a sender's type is the lower bound of the message sent is self-fulfilling.

a tradeoff. There is more frequent disclosure in the flexible language regime and more precise disclosure in the precise language regime. Sophisticated receivers, who are not fooled by vagueness, form more accurate beliefs under flexible than under precise language. Naive receivers form more accurate beliefs under precise language. Importantly, information transmission, i.e. the average accuracy of receivers' beliefs, is higher under precise language, irrespective of the proportion of naive receivers.

The Experiment. The experiment compares a Flexible and a Precise treatment that reflect the distinction between the two language regimes in the model. In both treatments, a sender's type is uniformly distributed over the integers from 0 to 5. A sender in the Flexible treatment can disclose any interval containing her actual type. For example, a sender with type 2 could disclose that her type belongs to the interval between 2 and 5. A sender in the Precise treatment can only disclose her exact type or nothing.

The theoretical predictions are borne out in the experimental data. Many senders are apt spin doctors. In Flexible, they use vague messages and the exact form of the modal message we observe is remarkably close to the one predicted by the model. In Precise, sender behavior reflects a threshold equilibrium in which only high types disclose. Senders disclose more in Flexible than in Precise and only a minority of senders in both treatments does not behave according to the theoretical predictions.

Validating the model's key assumption, we find evidence for the existence of two distinct receiver types, i.e. naives and sophisticates. We categorize receivers as either sophisticated or naive on the basis of their guesses and find that the average naive receiver makes smaller mistakes in Precise than in Flexible. Instead, depending on the specification, the average sophisticated receiver makes larger or equally large mistakes in Precise.

In encounters with rational senders, information transmission is significantly higher in Precise. When we consider all senders, the treatment effect of precise language on information transmission is positive but statistically insignificant because a very small number of observations in Precise feature a sender making the outlier mistake of not disclosing the highest type.

Policy implications. The exploitation of naive receivers may not be the only rationale for the use of vagueness in all field settings. For example, a sender may resort to vagueness to communicate that there is uncertainty about her precise type, but not about her belonging to some broader category. For this reason, it is crucial that the experimental laboratory allows us to strip the decision-making environment of any confounding drivers of vagueness and focus on its use in the exploitation of receiver naivite. Moreover, in contrast to most field settings, we can exogenously vary the language at a sender's disposal and characterize the exact nature of receivers' misinference. Our treatment comparison can then shed light on how policies that impose precise language on senders affect information transmission via their differential impact on sophisticated and naive receivers' inference.

Such policies are possible for all of our motivating examples and often feasible where mandatory disclosure is not (see section 3.5 for a discussion). In some settings, they are already in place. For example, Germany's main certifier of consumer products, Stiftung Warentest, gives products and services a precise mark and a vague summary category like "very good". It imposes precise language by legally requiring disclosures of its certification to contain the precise product rating. Similarly, the NGO Consumer Reports in the United States allows firms to post links to their online reports using neutral language, but does not allow them to excerpt content selectively. Our results suggest that restricting flexibility improves average information transmission and redistributes rents to naive consumers.

The imposition of precise disclosure is also taking root in science. For example, publishing guidelines by the American Psychological Association require authors to disclose the exact p-value, effect size, degrees of freedom, and statistical test underlying a given result. Similarly, the field of economics has undergone a general move toward precise disclosure. For instance, authors of experimental studies increasingly commit to the exact specifications of statistical tests in pre-analysis plans and thereby, among other things, reduce the subsequent flexibility in presenting their research findings.

The self-incrimination clause of the fifth amendment of the United States constitution affords defendants the right not to testify against themselves in criminal cases.[2] In settings in which lying is impossible (because testimony has to be backed up by hard evidence) or undesirable (because the expected penalty of perjury exceeds its benefits) the defendant's choice between testifying in her own trial or "pleading the fifth" constitutes a voluntary disclosure game with the jury. Moreover, a majority of US courts take the position that voluntarily waiving the right against self-incrimination opens a defendant up to cross-examination on all issues relevant to the trial.[3] The right not to self-incriminate therefore imposes precise voluntary disclosure. In an influential court ruling, the majority opinion argues against allowing the defendant to "decide how far he will disclose what he has chosen to tell in part [...]" because "it must be conceded that the privilege is to suppress the truth, but that does not mean that it is a privilege to garble it."[4] Our results highlight

[2] See Amar and Lettow (1995) for a critical discussion of this privilege.
[3] See Yale Law Journal (1952) and Stanford Law Review (1962) for discussions of the waiver and how it has and should be interpreted by the law.
[4] See the opinion by judge Hand in United States v. St. Pierre, 132 F.2d 837 (2d circuit 1942).

a key distinction between suppressing the truth and garbling it by means of partial or vague disclosure and speak to the wisdom in prohibiting the latter.

In the next section we discuss our relation to the literature, before presenting the model in section 3.3. Section 3.4 describes experimental design and results and section 3.5 discusses policy implications.

3.2 Related Literature

The Precise treatment is based on experiments by King and Wallin (1991) and Jin et al. (2018). Our empirical contribution lies in being the first to compare the effect of Precise and Flexible language on senders' communication strategies, receivers' inference and information transmission. To better interpret our results, our model extends theory by Eyster and Rabin (2005) and Hagenbach and Koessler (2017) to derive the welfare implications of the different language regimes. Taken together, theory and treatment comparison allow us to speak directly to the merits of regulating vagueness in disclosure. Moreover, our theory-guided design allows us to be precise about the nature of the bias in receivers' inference we uncover and to test for several of its implications. These features are likely to improve the robustness of our findings as well as their portability to other settings.

The exploitation of flexibility in voluntary information disclosure has been documented for car sellers describing their cars on ebay (Lewis, 2011), business schools referring to third-party rankings (Luca and Smith, 2015), and researchers presenting their findings (Krawczyk, 2015; Brodeur et al., 2016). Relatedly, there is evidence that firms shroud (Brown et al., 2010), obfuscate (Ellison and Ellison, 2009; Ferman, 2015), or complexify (Ru and Schoar, 2016) unfavorable information about their products. In markets where voluntary disclosure is necessarily precise, nondisclosure often ensues. For example, producers of salad dressings do not voluntarily disclose fat content if it is high (Mathios, 2000), poor health maintenance organizations do not obtain independent accreditations (Jin, 2005), and movie studios avoid pre-release screenings to critics if a movie's quality is low (Brown et al., 2012, 2013). However, data limitations in the field have thus far kept researchers from studying the causal impact of different language regimes on information transmission and from characterizing the exact nature of receivers' misinference.[5] On the theoretical front, a series of papers following Gabaix and Laibson (2006) investigate the circumstances under which firms fail to educate their own and other firms' consumers about unfavorable product attributes or add-on costs. However, the role of the flexibility of language in firms' communication with 'behavioral' consumers has been largely

[5] See Dranove and Jin (2010) for a review of the theory and empirics of disclosure in economic applications and Loewenstein et al. (2014) for the psychological subtleties surrounding the analysis of disclosure games.

neglected.

Our paper contributes to a small literature that studies information disclosure in the experimental laboratory. While no previous experiment features the treatment comparison and theory-guided design that allows us to speak about the consequences of intervening in the language at senders' disposal, our precise language treatment follows Jin et al. (2018), who provide evidence for both incomplete unraveling and receiver naivete. Earlier studies by Forsythe et al. (1989), King and Wallin (1991) and Dickhaut et al. (2003) find evidence for full unraveling after a sufficiently high number of repetitions, albeit in a setting that features several receivers and auctioning mechanisms that potentially permit other explanations for players' behavior (Jin et al., 2018).[6]

Three contemporaneous experiments complement our findings in the flexible language treatment. Jin et al. (2019) study a mandatory disclosure game in which senders can complexify their disclosure by revealing their type as the sum of a string of numbers. They find that low sender types make use of complexity and that some receivers are fooled by it because they are overconfident in their ability to interpret complex messages. In contrast, our results suggest that a lack of strategic sophistication leads to the loss in information transmission associated with vague disclosure. Hagenbach and Perez-Richet (2018) conduct an experiment that allows for vague messages. Instead of varying the language at a sender's disposal, they vary the sender's incentive structure. Like us, they find that types who wish to be perceived as another type are more likely to use partial or nondisclosure. They also find that receivers are better off under acyclical incentive structures, i.e. games in which masquerading incentives are not circular. Li and Schipper (2020) study a voluntary disclosure game in which senders can disclose any set of types that contains their actual type. Their experiment does not feature a treatment comparison of imposed precision and their version of flexibility. Still, senders use vagueness in a way that is reminiscent of senders' strategy in our experiment.

Strategic information revelation is also often analyzed within the cheap talk framework (Crawford and Sobel, 1982). In cheap talk games the sender is unconstrained in her choice of messages and it is typically assumed that the sender's and the receiver's interests are at least partially aligned. Even when all receivers are perfectly rational, full information revelation does not obtain in equilibrium, i.e., communication is inherently vague. Moreover, while the presence of naive receivers may induce the sender to deceptively inflate her messages, it may also enhance type separation and overall information transmission

[6] Also see Benndorf et al. (2015) for an unraveling failure that is driven by senders' bounded rationality and Brown and Fragiadakis (2018) for a receiver misinference that is not based on a lack of strategic sophistication.

(Ottaviani and Squintani, 2006; Kartik et al., 2007; Chen, 2011).[7] In the experimental literature, the benefits of vagueness in cheap talk communication have been shown to arise in various settings. In a public good game, Serra-Garcia et al. (2011) find that vague communication can be socially valuable when truthful communication conflicts with efficiency. In an coordination game, Agranov and Schotter (2012) show that a benevolent announcer may resort to vague announcements of payoff states to facilitate coordination. In an otherwise canonical cheap talk game, Wood (2016) documents that information transmission is higher when the message space consists of coarse rather than precise messages.

3.3 The Model

3.3.1 Setup

A sender (S) and a receiver (R) play a persuasion game in which R wants her guess to be as accurate as possible, while S aims at maximizing R's guess. When the state of nature is ω and R's guess is g, then S's payoff is $U_S = g$. R's payoff is $U_R = -(\omega - g)^2$, which implies that she finds it optimal to guess her expectation of the state. At the initial stage, ω is drawn from a continuous uniform distribution with support $\Omega = [0, 1]$.[8] S privately observes her type ω and sends a message m before R makes a guess.

Since S cannot make false statements, her message must always include her true type. Beyond this common requirement, we consider two alternative communication regimes, which we refer to as precise and flexible language. Under precise language, the set of messages available to type ω is $\{\omega, \Omega\}$, i.e. S can either reveal her type exactly ($m = \omega$) or remain silent ($m = \Omega$). Under flexible language, the set of messages available to type ω is the union of all convex and closed subsets of Ω containing ω, which also includes the option to remain silent.[9] While we represent the choice to remain silent with the coarsest message (i.e. $m = \Omega$) for ease of notation, in our interpretation this is conceptually distinct from actively disclosing as in the case of other messages. Thus, we say that a given type discloses only when she sends a message that conveys at least some information

[7] For similar reasons, information transmission may increase when messages get more vague due to an external noise that perturbs sender's communication (for example, see Board et al. 2007).

[8] Appendix C.2.2 considers non-uniform priors.

[9] An alternative specification of flexibility might allow senders to disclose any set of types that includes their actual type rather than constraining them to disclose an interval. We favor our modelling strategy for two reasons. First, disclosure of a disconnected set of types is rarely observed in the field, presumably because it is less natural and would tip naive receivers off. Modelling this tipping off explicitly would require that receivers' degree of naivite depends on the message, an unnecessary complication of the exposition. Second, imposing that disclosure has to take the form of an interval in the experiment helps to minimize the difference in complexity between precise and flexible language and thereby facilitates the interpretation of our treatment comparison.

even if interpreted at face value. Likewise, we measure the amount of disclosure as the probability that S sends a message other than $m = \Omega$, i.e. any strict subset of Ω.

Key to the analysis is that R may lack strategic sophistication. In particular, we take the posterior distribution of a fully naive R to coincide with the prior truncated over types for which the message sent is available. Upon message $m = \Omega$, the posterior of a fully naive R hence coincides with the prior, since that message is available to all types. Upon message $m = \omega$, her posterior is degenerate at ω. And upon message $m = [a, b] \subset \Omega$ with $a < b$, her posterior is uniform on $[a, b]$. Receivers' insufficient skepticism may hence stem from a failure to take into account the dependence of a sender's strategy on her type, in the spirit of cursed equilibrium (Eyster and Rabin, 2005).

We assume heterogeneity in sophistication as in Milgrom and Roberts (1986) and Hagenbach and Koessler (2017). In particular, we suppose that R is fully naive with probability $\chi \in (0, 1)$ and fully sophisticated with complementary probability. In Appendix C.2.1 we generalize the model by allowing for partially naive types and an arbitrary distribution of sophistication in the population.

A pure strategy of S specifies a message $m(\omega)$ based on her type. A pure strategy of R specifies a guess for a sophisticated and a naive type, which we denote respectively by $g(m)$ and $g_\chi(m)$, based on S's message. Our solution concept is a natural adaptation of perfect Bayesian equilibrium which takes into account that S's message is hard evidence and that R may not be fully strategic. In addition to the usual requirements, upon any off-the-equilibrium-path message, the support of R's posterior should not include any type for which that message is unavailable. Moreover, the guess of a naive R must be optimal given her possibly wrong and at least partially exogenously given beliefs. We restrict our attention to pure strategy equilibria and adopt the convention that S refrains from disclosing whenever indifferent.[10]

3.3.2 Predictions

When language is flexible, S elects to disclose an interval that spans her type and the highest type. As the equilibrium is necessarily fully separating,[11] this strategy is optimal in that it maximally inflates the guesses of a naive R.

[10]Given R's payoff function, her optimal action is always unique. As for S, the set of types that can have multiple optimal actions in equilibrium has zero measure.

[11]The reason for why full separation necessarily obtains is that, for any candidate equilibrium pooling message, the highest type in the pool always has access to another message that would strictly raise the guess of both a sophisticate and a naive R.

Proposition 1 (Equilibrium under flexible language)**.** *Under flexible language, in equilibrium* $m(\omega) = [\omega, 1]$, $g\big([a,b]\big) = a$ *and* $g_\chi\big([a,b]\big) = (a+b)/2$.[12]

Proof. See Milgrom and Roberts (1986). □

When language is precise, S finds it optimal to disclose if and only if her type is sufficiently high. The marginal sender is indifferent between perfectly revealing her type and remaining silent, which induces a higher guess from a naive R (the prior mean) but a lower guess from a sophisticated R (the average silent type). Then, higher types indeed find it optimal to disclose and lower types to remain silent. Also, the disclosure cutoff is lower than the prior mean and increases with the proportion of naives.

Proposition 2 (Equilibrium under precise language)**.** *Under precise language, there exists a unique cutoff* $\omega^* = \frac{\chi}{1+\chi}$ *such that in equilibrium:*

$$m(\omega) = \begin{cases} \omega & \text{if } \omega > \omega^* \\ \Omega & \text{if } \omega \leq \omega^* \end{cases} \quad g(m) = \begin{cases} \omega \text{ if } m = \omega \\ \frac{\omega^*}{2} \text{ if } m = \Omega \end{cases}; \quad g_\chi(m) = \begin{cases} \omega \text{ if } m = \omega \\ \frac{1}{2} \text{ if } m = \Omega \end{cases}.$$

Proof. See Eyster and Rabin (2005).[13] □

Comparing the two propositions, we can derive the following predictions about differences in players' behavior between the flexible and precise language regime.

Proposition 3 (Predictions on differences in behavior)**.** *For any given* $\chi \in (0,1)$

1. *Sender behavior:*
   - (a) *there is more disclosure under flexible than under precise language;*
   - (b) *the average disclosing type is higher under precise than under flexible language.*
2. *Receiver behavior:*
   - (a) *R's expected guess is lower under precise than under flexible language and in both cases it exceeds S's expected type;*
   - (b) *under both precise and flexible language R's expected guess increases with χ.*

[12] The equilibrium is unique up to a multiplicity in beliefs and guesses of a sophisticated R upon off-the-equilibrium-path messages. For ease of exposition, we are adopting the convention that these beliefs take the same form as in the case of on-the-equilibrium-path messages. Any other belief and associated sequentially rational guess which is sufficiently skeptical to deter S from sending a message $[a,b]$ with $b < 1$ is also part of an equilibrium.

[13] In Eyster and Rabin (2005), the disclosure cutoff is $\omega^* = \frac{1}{1+\chi}$ and it is types $\omega < \omega^*$ who disclose, since S's payoff is decreasing in R's guess. Also, as detailed in Appendix C.2.1, R's naivete takes a slightly different form.

Proof. See section C.1.1 in the appendix. □

Predictions on senders' behavior hinge on the fact that all types disclose under flexible language while only sufficiently high types disclose under precise language. Predictions on receivers' behavior are driven by the guesses of a naive R, since, given the Bayesian consistency of rational beliefs, the average guess of a sophisticated R always coincides with the prior mean. For any realization of S's type, the guess of a naive R is higher under flexible than under precise language, which explains the first part of prediction 2a, and always equal or higher than S's type, which explains the second part. The average guess of a naive is therefore higher than that of a sophisticate, which entirely drives prediction 2b under flexible language. Under precise language, it is also the case that the average guess of a naive increases with χ since, as S discloses less often, the set of S's types she overestimates increases.

To consider players' welfare, recall that the expected utility of S is simply the guess she expects to induce in R, while the expected utility of R is the accuracy of her guess measured by the mean squared error.[14] For a sophisticated R, this error boils down to the expected residual variance upon S's disclosure. For a naive R, the error also incorporates the systematic bias in her updating that S's strategy introduces. Throughout, we will use the terms R's expected payoff and information transmission interchangeably. Normatively, we take the stance that a social planer is interested in maximizing information transmission.[15] Moreover, the ex-ante and ex-post qualifications refer respectively to whether the expectation is computed unconditionally or conditionally on the player in question having observed her type (i.e., the state for S and the sophistication level for R).

Proposition 4 (Predictions on differences in payoffs). *For any given $\chi \in (0, 1)$*

3. *Sender payoff:*

   (a) *the ex-ante expected payoff of S is higher under flexible language than under precise language;*

[14] While R's preference ranking over flexible and precise language can only be defined with respect to a specific loss function in our setting, our results are robust to the use of the mean absolute error, i.e., $\mathbb{E}|g - \omega|$. When we present our experimental data we use this alternative measure, which is more directly interpretable in that it assigns no heavier relative penalty to larger errors.

[15] Our main rationale for focusing on information transmission and for not also taking into account the sender's payoffs in our discussion of policy is that any surplus the sender obtains relative to the full-rationality benchmark derives entirely from deception as opposed to an underlying economic fundamental. In common applications of persuasion games, such as sales or financial disclosures by managers, the sender's payoff can be thought of as a price or salary and therefore constitutes a pure transfer. On a practical note, results on the sum of sender and receiver payoffs would be sensitive to the exact scaling of players' payoffs, which would introduce an element of arbitrariness.

(b) *the ex-post expected payoff of S is higher under flexible language than under precise language if and only if S's type is not too low (in particular, it is true for all types who disclose under precise language).*

4. *Receiver payoff:*

(a) *the ex-ante expected payoff of R is higher under precise than under flexible language;*

(b) *the expected payoff of a sophisticated R is higher under flexible language and the expected payoff of a naive R is higher under precise language.*

Proof. See section C.1.2 in the appendix. □

Ex ante, S always prefers flexible language because it allows her to more strongly inflate the expectation of a naive R and it also offers more opportunities to do so. In spite of this, sufficiently low types still prefer precise language ex post, since it allows them to pool with higher types even in the eyes of a sophisticate. By the same token, a naive R prefers precise language, since it limits the scope for deceiving her, while a sophisticated R prefers flexible language, since it allows her to always perfectly infer the state.

The difference in the preference of sophisticates and naives for precise and flexible language has been noted by Hagenbach and Koessler (2017).[16] In the appendix we show that this result, obtained in a setting of two extreme levels of sophistication, naturally extends to settings where sophistication varies continuously in the population (proposition 5). More importantly, our setting allows us to sign the overall effect of language on R's payoffs: R prefers precise language. Intuitively, since a naive R is deceived more frequently and more severely under flexible language, the resulting loss has a substantially larger magnitude than the loss of both a naive and a sophisticate under precise language. Given the opposing language preferences of naives and sophisticates, one may think that R's ex-ante welfare could be higher under flexible language if the population of receivers is mostly sophisticated. This is not the case because the presence of more sophisticates also disciplines S's disclosure behavior under precise language (i.e., as χ goes to zero, so does the disclosure cutoff ω^*), which fosters information transmission to both sophisticates and naives.

In Appendix C.2.1, we show that R's preference for precise language is robust to arbitrary distributions of naivete in the population. In Appendix C.2.2, we use numerical simulations to demonstrate that it holds for a large class of non-uniform priors over the

[16]Our precise and flexible regimes correspond respectively to simple and rich language in the terminology of Hagenbach and Koessler (2017).

state. There, we also provide some quantitative measure of R's welfare gains from precise language and build some more intuition behind this result by identifying the features of prior distributions that can give rise to rare counterexamples.

3.4 The Experiment

3.4.1 Design

The experiment was programmed in zTree (Fischbacher, 2007). A total of 158 subjects participated in 8 sessions at the Munich Experimental Laboratory for Economic and Social Sciences (MELESSA) in the spring of 2017.[17] One session lasted for about 45 minutes and the average earnings (including a €4 show-up fee) were €15.05, with minimum earnings of €5.90 and maximum earnings of €23.50. The instructions were read aloud by the experimenter. Screenshots of the decision screens are gathered in Appendix C.6 and instructions and payoff tables can be found in Appendix C.7.

The experiment featured a between subject design that compared two variants of a disclosure game. At the beginning of the experiment, subjects in both treatments were randomly assigned to the role of a sender or the role of a receiver. A subject remained in her assigned role for the duration of the experiment. All subjects played 15 rounds of the disclosure game. In each round, a subject played the game with a randomly selected anonymous partner in the opposite role.

It was common knowledge that a sender's type ω was drawn in each round from the set $\{0, 1, 2, 3, 4, 5\}$ and that each type was equally likely. After privately observing her type, a sender decided on a message to send to the receiver. Our two treatments differed only in the type of messages senders were able to send.

In the Flexible language treatment (80 subjects), the sender was allowed to send any interval containing her type.[18] In the Precise language treatment (78 subjects), the sender could either disclose her precise type or do not disclose. In Flexible, senders were therefore able to send vague messages and while any feasible message in Precise was also feasible in Flexible, the reverse was not true. In the case of nondisclosure, the receiver was notified that "the sender did not send a message" in both treatments. Figure 3.1 depicts two messages a sender of type 2 might send in the different treatments.

[17] We piloted our design with 58 subjects in the winter of 2016. Here we organized the treatment variation in a within-subject fashion and find similar results.

[18] While sending an interval that contains all possible types was not allowed, the equivalent strategy of nondisclosure was always at a sender's disposal.

(a) Precise Message **(b)** Vague Message

Figure 3.1 Examples of Messages

After seeing the sender's message, the receiver stated her guess about the sender's type, i.e., $g \in \{0, 0.5, \ldots, 4.5, 5\}$. The receiver's action space was coarsened so that both sender and receiver payoffs could be communicated in the form of digestible payoff tables rather than relying on subjects calculating their payoffs by themselves. While the sender was incentivized to induce the highest possible guess in the receiver, the receiver was paid for accuracy. Subjects were paid in probability points and for a single randomly selected round. After each round, subjects received information about the receiver's guess, the sender's type and the probability points they earned.

A receiver's points depended on her guess and the sender's type as follows

$$p_R = \frac{110 - 20|\frac{\omega - g}{1.37}|^{1.4}}{110}$$

A sender's points depended only on the receiver's guess:

$$p_S = \frac{110 - 20|\frac{5 - g}{1.37}|^{1.4}}{110}$$

The probability points p a subject earned in the payoff-relevant round then determined the likelihood of winning a €8 prize. For example, a subject in the receiver role was paid according to a lottery that yielded a relatively high prize of €8 with probability p_R and a lower prize of €1 with the complementary probability $1 - p_R$. Paying subjects in probability points makes them less liable to the influence of risk preferences (Roth and Malouf, 1979; Hossain and Okui, 2013; Harrison et al., 2014). To make sure that subjects understood the incentive structure we provided them with payoff tables that mapped any constellation of receiver guess and sender type into the relevant probability points and let them solve comprehension tasks before the experiment.

After the main part of the experiment, we elicited subjects' "out-of-sample" beliefs about behavior in the pilot experiment. Senders stated the distribution of receiver guesses upon nondisclosure and receivers stated their belief distribution over non-disclosed sender types.[19] Subjects were paid for being close to a variable's empirical distribution in the

[19]In additional unincentivized elicitations in Flexible, we asked senders about the average receiver guess in

pilot sessions (see Appendix C.7.4 for details). These elicitations facilitated a rationality check for senders (see Appendix C.5) and a consistency check for our naivite classification of receivers. Finally, a very short post-experiment survey collected some additional sociodemographic data.

3.4.2 Results

We first describe participants' behavior in the two treatments and then analyze information transmission. Our analysis is based on data that pools observations across rounds. Appendix C.3 provides results on how player behavior evolves over time. For all statistical tests we report p-values from a two-sided t-test that comes from a regression-based approach with robust standard errors clustered at the subject level.

Behavior

Flexible Language Treatment. According to the theory, a sender in the Flexible treatment discloses an interval that spans her type and the upper bound of the type space. Figure 3.2a depicts the average lower and upper bounds of the messages sent by different sender types. Observed messages are in line with the predictions of the model. Upper bounds are close to the highest type and lower bounds increase with the type. Modal messages, also depicted in the figure, almost perfectly coincide with the theory's predictions. The only exception is provided by senders of type 1, who remain silent more often than they send their predicted message.

As a first step toward analyzing receiver behavior, we normalize guesses. Given a guess g and a message with lower bound $\underline{\omega}$ and upper bound $\bar{\omega} > \underline{\omega}$, the normalized guess is

$$g_n = \frac{g - \underline{\omega}}{|\bar{\omega} - \underline{\omega}|}.$$

The normalization allows for the comparison of guesses induced by different messages. Normalized guesses range from 0 to 1 and are only defined for nondisclosure or vague disclosure. A fully naive normalized guess always takes a value of 0.5. The theoretical prediction for a sophisticated normalized guess in Flexible is 0 for all messages.

Figure 3.2b shows the distribution of normalized guesses. The bimodal distribution with mass points at 0 and 0.5 vindicates our model's assumption that there are two distinct

the pilot session after receiving the messages $\{1,2,3,4,5\}$, $\{2,3,4,5\}$, $\{3,4,5\}$, and $\{4,5\}$ and receivers about the most likely message of all six possible sender types.

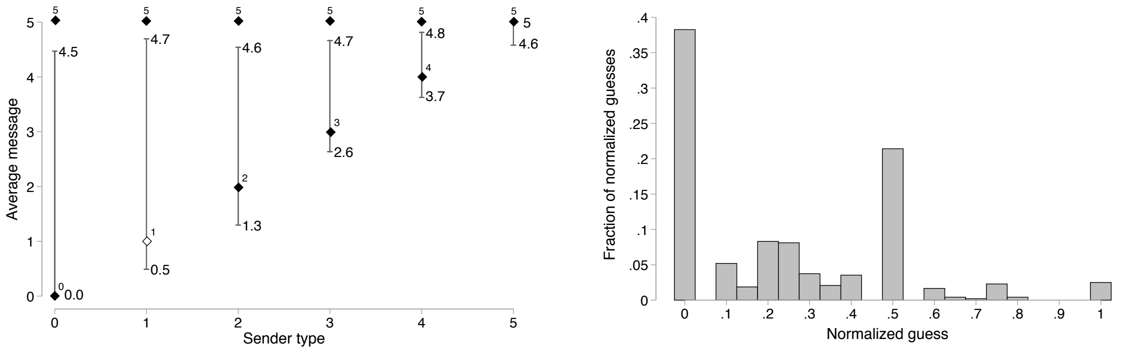

(a) Messages **(b)** Normalized Guesses

Figure 3.2 Behavior in the Flexible Treatment

Notes: (a) Solid lines show the avg. lower and upper bounds of all messages sent. Diamonds show the model's predicted lower and upper bounds; diamonds are black when predictions coincide with the modal message in the experiment and hollow if not. Average and modal messages include nondisclosure. (b) Bars show the distribution of normalized guesses.

receiver types: sophisticates and fully naive receivers.[20] We find that receivers' average belief upon observing nondisclosure or receiving a vague message is upwardly biased. While the average normalized guess is at about 0.25, senders' average normalized type is significantly lower at 0.13 (p-value < 0.001).[21] Instead, all receivers are able to rationally interpret singleton intervals, i.e., a precisely disclosed type.

Precise Language Treatment. In the presence of naive receivers, our model predicts that precise language will give rise to a threshold equilibrium with nondisclosure on behalf of low types and disclosure on behalf of types above the threshold. Figure 3.3a depicts disclosure rates by sender type. In line with an equilibrium threshold of around 2, the disclosure rate is almost zero for the lowest two types, 40% for type 2, and above 80% for the highest three types. Note that disclosure rates of less than 100% for the highest types imply a slight departure from our hypothesis of sender rationality.[22]

All receivers are able to rationally interpret a precisely disclosed type. Figure 3.3b depicts the distribution of normalized receiver guesses upon nondisclosure. We observe a bimodal distribution with mass points around 0.2 and at 0.5.[23] Because of the threshold strategy, the sophisticated guess upon nondisclosure is now larger than zero. In particular, a receiver's empirical best response is equal to the average non-disclosing type, whose normalized value is equal to 0.25. Therefore, the histogram's first mode reflects

[20] We can reject the null hypothesis of unimodality using the Dip Test introduced by Hartigan and Hartigan (1985) (p_{Dip}-value < 0.001).

[21] The normalized type is the sender counterpart of the normalized guess and is given by $\omega_n = \frac{\omega-\omega}{|\bar{\omega}-\omega|}$.

[22] However, after the initial five rounds the disclosure rate of high types increases markedly, e.g., for sender type 5, it increases from 70.8% to 92.9%.

[23] We can reject the null hypothesis of unimodality using a Dip Test (p_{Dip}-value < 0.001).

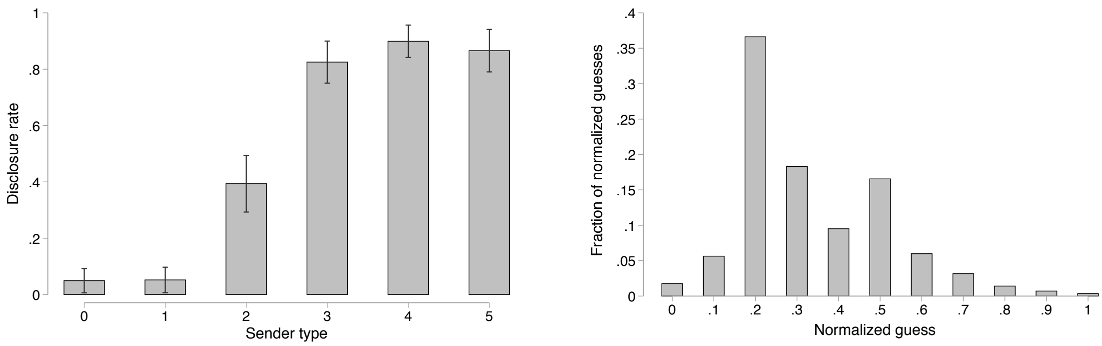

(a) Disclosure Rates

(b) Normalized Guesses

Figure 3.3 Behavior in the Precise Treatment

Notes: Graph (a) shows 95% confidence intervals around the avg. disclosure rates. Graph (b) shows the distribution of normalized guesses.

the accurate beliefs of sophisticated receivers, whereas the second mode corresponds to the beliefs of a fully naive receiver.

Receivers average normalized guess upon observing nondisclosure is 0.33, which reflects a significant overestimation of the average normalized non-disclosed type of 0.25 (p-value $= 0.007$).[24]

Treatment Comparison. Disclosure rates are higher in Flexible, where senders disclose 75% of the time, than in Precise, where they disclose 51.5% of the time (p-value < 0.001). This result seems to be driven by differences in disclosure strategies, as the average disclosing sender type is significantly higher in Precise than in Flexible (3.59 versus 3.08; p-value < 0.001).[25]

We observe that the average receiver guess is slightly lower under Precise than under Flexible (2.66 in Precise versus 2.85 in Flexible; p-value $= 0.095$). This difference suggests that, as the theory predicts, the average sender is better off in Flexible.

[24] Appendix C.5 depicts senders' out-of-sample predictions of the pilot's receivers' guesses upon non-disclosure. Results are reflective of high average sender rationality and an unbiased understanding of receiver behavior. Matching actual receiver behavior in the experiment, senders' predictions feature a modal normalized guess of 0 in Flexible and 0.2 in Precise as well as substantial weight on high, naive guesses. Therefore, the average sender appears to be best-responding to unbiased beliefs about receiver behavior.

[25] Although not significantly different, results also agree with the predicted direction for the differences in average non-disclosing types: 1.26 in Precise treatment versus 1.07 in Flexible language (p-value $= 0.310$).

Information Transmission

We measure information transmission by receivers' *mistakes*, which themselves are given by the absolute difference between a receiver's guess and a sender's type. Perfect information transmission corresponds to a mistake of zero.

Overall Information Transmission. Our model predicts that average receiver mistakes are lower in Precise. Table 3.1 shows the determinants of receivers' mistakes. Column (1) depicts an OLS regression of receiver mistakes on the treatment and tells us that the treatment effect of Precise on average mistakes is negative, but insignificant. The insignificance of the treatment effect is driven by the minority of sender choices that do not conform to our theoretical predictions. To test this hypothesis, we restrict the sample to the 828 observations that feature theory-conforming sender behavior. In Flexible, such behavior takes the form of a message that spans the sender's type and 5. In Precise, it takes the form of a threshold strategy, whereby only types of 2 or higher disclose. Here, the threshold of 2 is the best response to the distribution of receiver guesses upon nondisclosure.

Column (2) focuses on the 70% of interactions in which sender behavior conforms *exactly* to the theory. In these cases, restricting our senders to the use of Precise language leads to lower average receiver mistakes. The significant treatment effect emerges because in focusing on theory-conforming behavior, our data restriction eliminates a very small number of outlier observations driven by sender mistakes that disproportionately occurred in Precise.[26] Imposed precision therefore improves information transmission in the absence of pronounced sender irrationality that is unbalanced across language regimes.

A Typology of Players. The theory predicts that moving from Flexible to Precise decreases the average mistakes made by naive receivers and increases the average mistakes made by sophisticated receivers. A corollary of this prediction is that the interaction effect of imposed precision and a receiver's naivete on mistakes is negative. In order to test these predictions, we use our experimental data to classify receivers as naives and sophisticates.

A normalized guess is fully naive if it is equal 0.5. We arrive at our measure of individual receiver naivete by dividing the number of rounds in which the receiver stated a fully naive guess by the number of rounds in which the receiver did *not* encounter precise dis-

[26] In particular, the treatment effect on average information transmission is also significant at the 10% level if we merely drop the 12 observations (1% of total observations) that feature a sender of type 5 who does not disclose and thereby generates a disproportionately large outlier receiver mistake. 11 of these observations occurred in Precise.

Dep. variable:	(1) Mistake	(2) Mistake	(3) Mistake	(4) Mistake	(5) Mistake	(6) Mistake
Precise (d)	-0.0610	-0.159**	-0.252***	-0.231**	0.140*	0.0250
	(0.0676)	(0.0746)	(0.0754)	(0.0856)	(0.0787)	(0.0826)
Round	-0.0298***	-0.0207***	-0.0331***	-0.0342***	-0.0257***	-0.00932*
	(0.00582)	(0.00498)	(0.00995)	(0.00871)	(0.00674)	(0.00465)
Constant	1.647***	1.742***	2.168***	2.322***	1.259***	1.277***
	(0.106)	(0.120)	(0.165)	(0.162)	(0.115)	(0.130)
Type dummies	Yes	Yes	Yes	Yes	Yes	Yes
Incl. sender choices	All	Theory-conforming	All	Theory-conforming	All	Theory-conforming
Incl. receivers	All	All	Naives	Naives	Soph.	Soph.
R^2	0.172	0.456	0.331	0.605	0.135	0.454
Observations	1185	828	510	360	675	468

Table 3.1 Regressions of the Treatment Effect on Receivers' Absolute Mistakes

Notes: Robust standard errors clustered at the subject level in parentheses; $^{*}p < 0.10,^{**}\ p < 0.05,^{***}\ p < 0.01$.

closure. If this ratio is smaller than 0.15, then we say that a receiver is "hardly ever naive" or sophisticated. Otherwise, a receiver is deemed naive. Applying this classification procedure, we find that 57% of the receivers in our sample are sophisticated.[27]

Columns (3) and (4) of Table 3.1 repeat the regression models of columns (1) and (2), but include only naive receivers. Regardless of whether or not we only include theory-conforming sender behavior, the treatment effect on naive receivers is negative and significant. In column (5), we see that, as the theory predicts, the treatment effect on sophisticated receiver's mistakes is positive and weakly significant. However, the result is not robust to considering only theory-conforming sender behavior and therefore disappears in column (6).[28]

[27] The fraction of naive receivers is higher in Precise (61.5%) than in Flexible (52.5%). However, this difference is not robust to different classification criteria. In general, the proportion of naives is slightly higher in Precise if the classification is based on the frequency of fully naive choices (as in our main classification) and slightly lower if the classification is based on the proximity to best-response behavior, as some classifications in Appendix C.4. Therefore, there is no reason to suspect that the treatment effect on overall information transmission is driven by differences in the proportion of naives across treatments. To validate our naivite classification, we may ask whether receivers' out-of-sample beliefs, elicited after the experiment, about the type of non-disclosers in another experiment vary systematically according to their classification. Indeed, we find that the average normalized guess of naives is 1.28 whereas the average guess of sophisticates is 0.86 (p-value $= 0.022$).

[28] Results in Table 3.1 are largely robust to adding session clusters as well as to using the probability points earned by receivers as the outcome measure. When we cluster at the session level, results on information transmission to average receivers and naive receivers remain unaltered, while sophisticated receivers

In Appendix C.4 we repeat the regressions in columns (3) through (6) for several alternative classifications of naivete and sophistication, including subjects' high school math grade and various notions of empirical best response. In the majority of specifications, naive receivers make significantly smaller mistakes under precise language and sophisticated receivers make insignificantly larger mistakes under precise language. The appendix also shows that the data bears out the corollary of our model's predictions: for all classifications, moving from flexible to precise language leads to relatively smaller mistakes for naive receivers, i.e., the interaction effect between the precise treatment and naivete on mistakes is negative.

The negative coefficient of the variable *Round* in all six regressions indicates that there is a negative time trend in receiver mistakes. Appendix C.3 provides a more detailed analysis of the evolution of play on behalf of both senders and receivers. Subjects in both roles learn as rounds progress, but for the most part there is no time-varying treatment effect on information transmission.

3.5 Discussion and Conclusion

Our model and experimental data suggest that information transmission can be increased by restricting senders' flexibility in disclosing private information to receivers. Moreover, we find that a move to precise voluntary disclosure is likely to disproportionally benefit naive receivers. Since sophisticated receivers are (weakly) harmed by restricting flexibility, while naive receivers benefit, it is tempting to think that the effect of restricting flexibility on average receiver welfare is generally negative when there are many sophisticated receivers. However, this intuition is wrong: restricting flexibility improves information transmission for a broad class of distributions of strategic sophistication. When there are many sophisticates, precise language features (almost) full disclosure and still beats out the flexible language regime.

We have analyzed the disclosure game through the lens of sender rationality. In terms of the applications we have in mind, it is plausible that professional marketers are able to make cunning disclosure decisions and that high-paid attorneys are able to advise their clients on optimal disclosure strategies. And while senders and receivers are often drawn from the same population in the case of research, authors of papers naturally devote substantially more time and cognitive resources to a paper than a paper's readership is able to. Our theoretical results can accommodate and are robust to some sender irrational-

no longer fare better under flexible language. Using probability points further increases the pull of outlier observations because of the convexity of the payment schedule. As a result, the treatment effect on information transmission in column (2) is only significant at the 10% level.

ity. However, as our experiment shows, noisy behavior on behalf of senders can make it difficult to detect the benefits of precise disclosure in the experimental laboratory.

In our simple framework, an easy way to facilitate information transmissions is to legislate the mandatory disclosure of information. Where mandatory disclosure is feasible and unproblematic, our results suggest that it is crucial to also legislate precise language. However, for a number of reasons mandatory precise disclosure may often be infeasible or undesirable where the mere imposition of precision is not.

First, mandatory disclosure may be deemed unfair. Consider a policy maker's decision to regulate the disclosure of college rankings by colleges. While informative, rankings also contain an element of subjectivity and may be subject to dimensions, like students' entertainment facilities, that a college reasonably neglects. Therefore, it may be deemed unfair and invite resistance to force a college's disclosure of a given ranking. Nonetheless, conditional on a college's voluntary disclosure, imposing precision by prohibiting disclosure in selectively broad categories (e.g. "top 30") is likely to be less contentious. Concerns about fairness are also at the heart of arguments in favor of the self-incrimination clause of the fifth amendment.

Second, it may be prohibitively onerous for a regulator to determine whether a firm chose nondisclosure or simply lacked information. Consider a pharmaceutical company that tests one of its products only to find that the product has the unfortunate side effect of hair loss in 9 percent of the study's participants. In the case of nondisclosure, it may be hard for the regulator to find out whether a study was ever conducted. However, a press release that claims that "less than 10 percent" or "a small minority" of participants experienced hair loss could easily be flagged for vagueness.

Third, mandatory disclosure may yield perverse incentives. For example, consider a defendant's right not to self-incriminate. In its absence, law enforcement has an incentive to use coercion or even torture to extract an admission of guilt. In the case of markets, Matthews and Postlewaite (1985) and Polinsky and Shavell (2010) demonstrate that forcing firms to reveal their private information may ultimately hamper information transmission once firms' incentives to acquire information are taken into account.

The question of how the presence of naive receivers affects information transmission when senders are not exogenously endowed with private information about their type is an interesting avenue for future research. In particular, it is plausible that mandating precise language has a disincentive effect on information acquisition, given that it sets a limit on senders' ability to use information to deceive receivers. This would limit the benefits of imposing precision. At the same time, in other settings, flexible language may be even more harmful than our data suggests. Cain et al. (2005) show that the disclosure of a conflict of interest can lead advisors to give more biased advice by making them feel morally licensed to pursue their private goals. Because flexible language leads to both less information transmission and more disclosure (i.e., moral licensing), it may lead both to a greater underappreciation of an advisor's conflict of interest and to poorer advice.

Our results pertain to information transmission to an average receiver. But the ultimate desirability of precise language may hinge on the weight society attaches to different receiver types. For example, in the case of research, society may deem that information transmission to referees, who are mostly sophisticated, is initially more important than information transmission to the general public, who is more likely to be naive. Yet researchers may write up their findings in an attempt to persuade both of these audiences. It may then be the case that flexible language and its superior information transmission to sophisticated receivers ought to be favored. On a related note, in some settings vague messages may serve a more benevolent purpose than the exploitation of naive receivers. For example, an organisation or policy maker may resort to vagueness to communicate uncertainty about the exact type. Then, if precision were imposed, this would lead to overprecise beliefs on behalf of receivers.

Finally, while focusing on information transmission is sensible in our general setting, it is not always clear-cut how information transmission maps into welfare in specific applications. For instance, Ispano and Schwardmann (2018) show that when vertically differentiated firms compete for sophisticated and naive consumers through quality disclosure, inflated beliefs about low-quality products on behalf of naive consumers may improve welfare by exerting competitive pressure on the prices of high-quality products. Moreover, consumers might simply enjoy thinking of the Bordeaux region while drinking a blended wine from Roussillon.

Chapter 4

Complexity and Appropriation Interact in Affecting Compliance Behavior*

4.1 Introduction

Compliance decisions are often very complex, requiring that individuals process large amounts of information and rules, and file ample paperwork. Complexity of tax rules and the tax filing process are particularly cumbersome (e.g., Slemrod and Sorum, 1984; Benzarti, 2020), causing inattentive decision making (Abeler and Jaeger, 2015) as well as confusion (Feldman et al., 2016; Taubinsky and Rees-Jones, 2017). Tax payers in the Canadian province of Québec for example need to file up to 43 forms using an instruction guide of more than 100 pages (Vaillancourt et al., 2015). In the United States, the burdens associated with filing taxes have been estimated to cost about 1.2% ($200 billion) of the GDP (Benzarti, 2020).[1]

Governmental officials all over the world have recently started discussing the hypothesis that complexity contributes to the tax gap between tax that is owed and tax that is paid (Government Accountability Office, 2017; Luttmer and Singhal, 2014). While these official reports acknowledge that taxpayers may underclaim benefits, it is still believed that complexity triggers predominantly self-serving non-compliance, whether intended or not. In contrast, inattentive decision making or confusion are more likely to generate random deviations from required levels of compliance. To our knowledge there is no

*This chapter is based on joint work with Charles Bellemare and Florian Englmaier.

[1] See Slemrod and Sorum (1984) or Blumenthal and Slemrod (1992) for corresponding estimates obtained using survey data.

direct evidence of factors responsible for the hypothesis that complexity contributes to self-serving non-compliance.

In this chapter, we use a controlled experiment to investigate the effects of complexity on compliance behavior. We focus on the question whether complexity effects are linked to the appropriation of taxes (i.e, the purpose for which collected tax funds are spent). This allows us to study whether complexity induces random deviations from compliance rules or rather serves a selfish/prosocial purpose.

Subjects in our experiment first generate income in a real effort task before being randomly assigned into one of four treatments based on a 2×2 factorial design. This design varies complexity of compliance decisions and appropriation. In all treatments, subjects are asked to calculate the share of their generated income they should keep as take-home pay, with the residual share to be donated to a designated charity. We vary complexity by manipulating tax forms from the province of Québec (Canada). In Simple treatments, subjects are asked to calculate the share of their generated income they are required to keep by completing a single one page form requiring three data entries. In Complex treatments subjects are required to complete seven forms requiring 34 data entries. All forms (both in Simple and Complex) were calibrated such that subjects who make the correct calculations would be asked to keep exactly half of their generated income, with the remaining half to be donated to their designated charity. We vary the appropriation of donations by randomly assigning two different existing and certified charities across subjects. The first charity raises funds to facilitate stem-cell donations to newborns with blood cancer. The second charity is a luxury private yacht club located in Germany. Both organizations are certified as charitable organizations under German law, and hence donations to both are tax deductible. Yet, redistributing generated income to the yacht club plausibly triggers a stronger perception that donations to this charity as less morally justified. All subjects were asked to keep their share as take-home pay, leaving the remaining share in a closed envelope to be donated after the end of the experiment. There were no risks or penalties for non-compliance, ruling out these considerations from our analysis.

We find that complexity has no significant effect on compliance when taxes are distributed to a deserving charity. This mirrors results in Dwenger et al. (2016) who found no effects on compliance behavior when simplifying payment of Church taxes. Our results add to this and related findings suggesting a significant share of taxpayers are intrinsically motivated to comply with complex rules when taxes are distributed to a morally justified cause (Abeler et al., 2019). Conversely, complexity is found to have a significant effect on compliance when taxes are distributed to the luxury yacht club. We also find that, conditional on forms under Simple, subjects keep significantly more of their generated income when taxes are distributed to a (more) deserving charity. This effect is consistent with a pure morality effect suggesting how taxes are used matters for compliance. Related non-experimental evidence consistent with this finding is Torgler (2003) who finds that distrust in governments is positively associated with acceptability of tax evasion. Overall, we find a significant interaction effect between complexity and morality – non-compliance is significantly accentuated under Complex when taxes are distributed to the yacht club.

One explanation for the interaction effect is that complexity serves to reduce psychological costs from non-compliance. Alternatively, it might be the case that the appropriation of taxes affects the subjects' compliance motivation and hence makes them more sensitive to complexity variations. We provide further evidence showing that efforts to file the forms are largely independent from tax appropriation. Moreover, calculation mistakes on complex forms appear highly correlated with non-compliance when taxes are donated to the yacht club. This bias is not observable when the money is donated to the cancer charity; here, mistakes induce random deviations from the compliance rule. Both observations together are suggestive evidence that complexity facilitates self-serving non-compliance behavior.

The interaction of complexity and appropriation has implications in many areas. To start, officials designing tax policy need to take into account the perception of taxpayers concerning the efficiency and appositeness of government spending. When perceptions are favorable, taxpayers are able and willing to work through complex rules, offering a leverage for elaborate tax policy. Spiegler (2016) surveys a literature in behavioral industrial organization arguing that firms may profit from strategically introducing complex rules at the expense of customers who then have difficulties making correct value comparisons across market alternatives. Examples range from major industries such as insurance, retail banking, or telecommunications to the mundane task of supermarket shopping where the large variety of potential substitutes, nonlinear and frequently changing prices, and incommensurable measurement units complicate choices. This complexity can be explicit, for example, elaborate fee structures employed by retail banks, or long service contracts loaded with impenetrable jargon or implicit as the arcane reimbursement practices of insurance companies. While product complexity is hard to avoid in many cases, it is a common intuition that part of the complexity is in fact strategic, designed by firms to take advantage of consumers. Our results suggest such distributive unfair practices may affect compliance in these industries. Insurees, for example, may withhold or distort information when filing complex reimbursement claims, amplifying moral hazard problems in insurance markets.

The remainder of this chapter is structured as follows. Section 4.2 presents the experimental design used in the chapter. Section 4.3 presents our main results and discusses the policy implications. Section 4.4 concludes.

4.2 Experimental Design

At the beginning of an experimental session, each subject generated income by positioning sliders on their computer screens (Gill and Prowse, 2019). Each correctly positioned slider generated €0.40 for a subject. Subjects positioned the sliders over two rounds (120 seconds per round). In each round, the screen presented 48 sliders and was split in two sections "No. 1-24" and "No. 25-48". Subjects were informed that their final payout from the experiment would be monotonically increasing in the number of correctly positioned sliders, but they were not informed of the exact share of the generated income they would receive as a final payment for the experiment. They worked in isolation from one another at separate computer terminals.

After completing the slider task, all subjects privately received their generated income, written instructions, one envelope containing forms as well as two empty envelopes. Written instructions indicated that they had to calculate the share of their generated income they could keep as payment for the experiment using the forms supplied.[2] They were further instructed to place their shares in the empty envelope labelled "your share", and the remaining shares in the second empty envelope labelled "remaining share". The instructions also indicated that the content of the "remaining share"-envelopes would be donated to a designated charity.[3] Subjects were informed that researchers would only collect the envelopes labelled "remaining share" after everybody left the room, and that the content of these envelopes would be transferred to the designated charity. Subjects were also instructed to leave behind the forms filled to calculate the respective shares. Forms and content of "remaining share"-envelopes thus contained information to measure rule non-compliance as well as possible calculation mistakes. Risk aversion is ruled out in the above design as subjects faced neither probabilisitc audits nor penalties for non-compliance or mistakes.[4] As a result, form calculations and content of the "remaining share"-envelopes need not match. Moreover, no binding time restrictions were placed on subjects to complete the experiment. Subjects were nevertheless presented a reference

[2]The full instructions, translated from German, can be found in Appendix D.1.

[3]An overview of the number of correctly positioned sliders in both rounds was displayed on the computer screen. Screenshots can be found in Appendix D.2. Each workplace was equipped with a pen, sticky tape, and a calculator. The sticky tape was used to seal all envelopes after shares were allocated.

[4]We made sure that this was clear to subjects by allowing them to put the "your share"-envelopes in their bags and by telling them that these envelopes must not be opened.

time of 900 seconds which was reached by about 5% of subjects. Subjects left the laboratory after answering some socio-economic questions.

The experiment is based on a 2x2 between-subject factorial design, interacting complexity and the appropriation of the taxed money. Forms were either Simple (a one page form with three items to fill) or Complex (seven forms with a total of 34 items to fill). Under Complex, forms also incorporated *if*-conditions and required subjects to transfer intermediate calculations across the different forms. We utilize abstract formats of the tax forms used in the Canadian province of Québec to operationalize complexity.[5] Forms (under Simple or Complex) were calibrated for the experiment such that subjects who comply and make no calculation mistakes were asked to keep 50% of their generated income as payment for the experiment, with the remaining income to be placed in the "remaining share"-envelope for later distribution to the designated charity. The second treatment dimension varies the designated charitable organization the taxed money is donated to. Half of the subjects were informed that the content of the "remaining share"-envelope would be donated to the *Deutsche Knochenmarkspende* (in English: German Bone Marrow Donation Registry; hereafter DKMS). The other half of subjects were told contents would be donated to the *Bayrischer Yachtclub* (in English: Bavarian Yacht Club; hereafter BYC). Instructions for subjects presented the mission statement of each organization translated below:

DKMS

Wir besiegen Blutkrebs.

> "*The main activity of the DKMS is to improve the healing potential of leukemia and other life-threatening diseases of the blood-forming system by supporting bone marrow donations. One major part of DKMS is the DKMS umbilical cord blood bank, which collects, processes, stores, and mediates umbilical cord blood stem cell donations for newborns. (Information from www.dkms.de)*"

[5] Both a screenshot of the original Québec tax forms and the experimental versions can be found in Appendix D.3.

Bayerischer Yacht-Club

"*The main activity of the BYC is to professionally promote sailing with all its modern features and high standards. In addition, the social life outside the gates of Munich is cultivated. The BYC also has an exquisite restaurant in its Club-casino at the Lake Starnberg. (Information from www.byc.de)*"

DKMS and the BYC are both classified as charitable organizations ("gemeinnützig") under German tax law, making them eligible for tax-preferred donations. While both organizations are legitimate recipients of donations, donating to DKMS appeals to higher moral standards, while donations to BYC, an elite organization in Germany, is intended to invoke the idea of ineffective or wasteful spending and rather low moral standards. In the following section, Waste will denote treatment specific donations to BYC, while Moral will denote donations to DKMS. Similarly, Simple and Complex will denote treatment specific form complexity described above.

The experiment was programmed using zTree (Fischbacher, 2007). 320 subjects (80 per treatment cell) were recruited with the help of ORSEE (Greiner, 2015) and participated in 32 sessions of our experiment at the Munich Experimental Laboratory for Economic and Social Sciences (MELESSA) in the summer of 2017. Every session was supervised by the same experimenter. The core socioeconomic variables are balanced across treatments, suggesting successful treatment specific randomization; see Appendix D.4.

4.3 Results

4.3.1 Data

Subjects on average generated €16.86 of income in the slider task (3.88 std. dev., minimum of €0, maximum of €28). The empirical distribution of generated income is similar to what has been reported in other experiments using the same slider task (e.g., Gill and Prowse, 2019; Abeler and Jaeger, 2015). Figure 4.1 presents the distributions of generated income for each of the four treatment groups. Average earned income across treatments are similar, ranging from €16.45 in Waste/Complex to €17.36 in Moral/Complex.

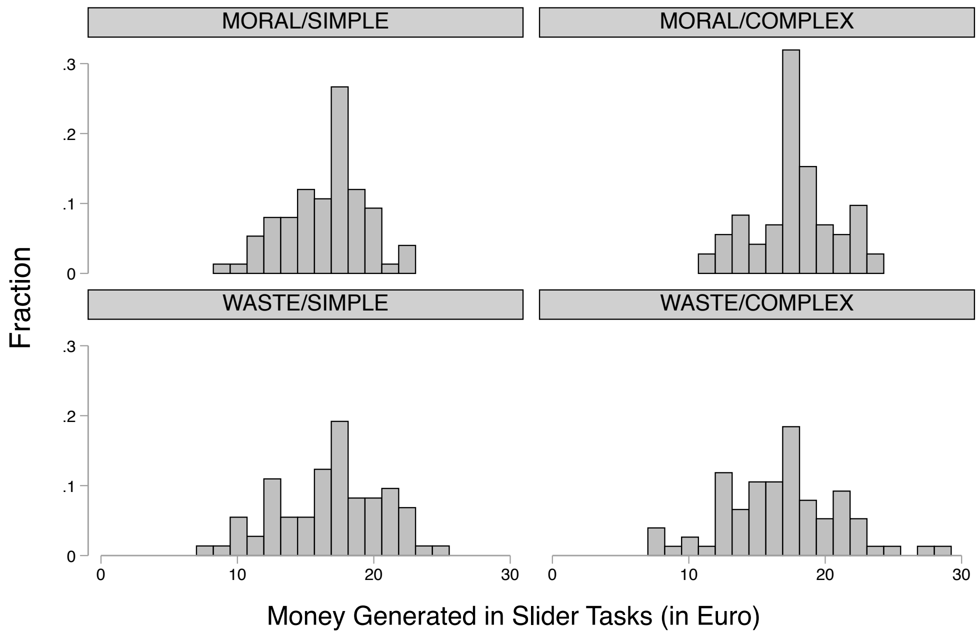

Figure 4.1 Distribution of Income from Slider Tasks

Distributions are not statistically different, with no pairwise two sample Kolmogorov-Smirnov tests rejecting the null hypothesis at usual significance levels (lowest p-value = 0.172). This is reassuring because subjects had no treatment-specific information available when generating their income.

296 of the 320 subjects completed all steps of the experiment and hence form our sample of analysis. Of the 24 subjects that are excluded from the analysis, two did not position a single slider correctly, 9 took home the forms, and another 13 left almost all items of the forms empty. These behaviors are not treatment-specific, reflected in the fact that our main results are similar when accounting for selective filing in a Heckman selection model (when appropriate) as shown in Appendix D.5. The net number of subjects per treatments are 75 in Moral/Simple, 72 in Moral/Complex, 73 in Waste/Simple, and 76 in Waste/Complex.

4.3.2 Compliance Behavior

Figure 4.2 presents the compliance behavior on the extensive margin. Compliers, over-providers, and evaders are defined as subjects who respectively donate 50%, more than 50%, or less than 50% of their generated income to their designated charity. We find that the proportion of compliers is significantly higher in Moral relative to Waste treatments (χ^2; $p < 0.001$). Pooling over morality dimensions, complexity has a negative effect on

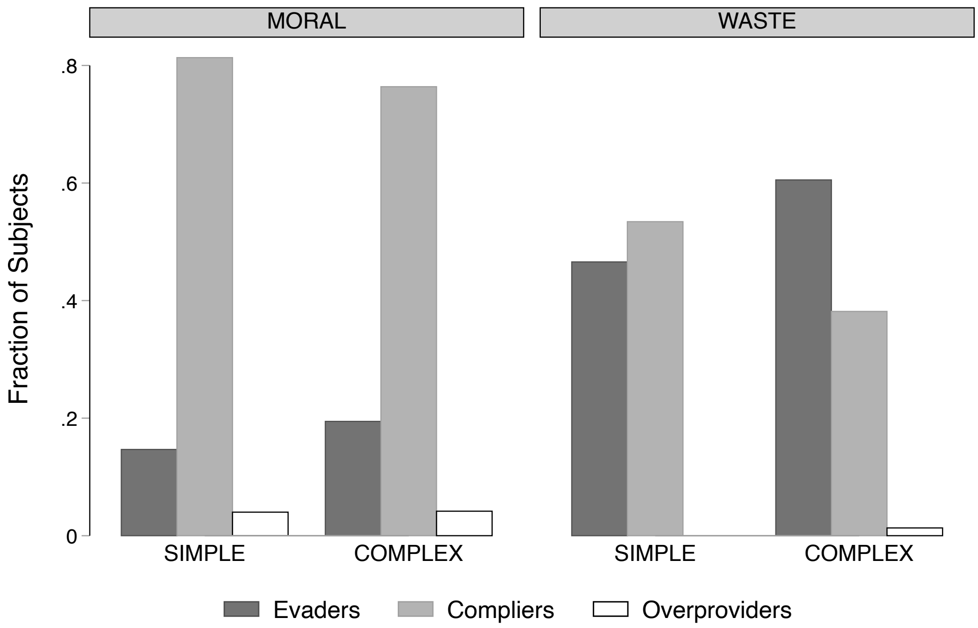

Figure 4.2 Compliance Behavior on the Extensive Margin

Notes: Compliers, overproviders, and evaders donate respectively 50%, more than 50%, or less than 50% of their generated income to the designated charitable organization.

the number of compliers (χ^2; $p = 0.055$).[6] Testing for complexity effects separately, we observe significantly fewer compliers due to complexity under Waste but not under Moral (χ^2; $p_{\text{Waste}} = 0.061$; $p_{\text{Moral}} = 0.463$). These results suggest that subjects are willing to comply and work through form complexity when the designated charitable organization is morally justified. Finally, we observe 4% of subjects being overproviders under Moral irrespective of form complexity, reflecting that prosocial subjects are not bound to limit their donations to the rule set in the experiment.

Treatment effects on the intensive margin are presented in Figure 4.3. All graphs plot the corresponding distribution of donations per treatment along with sample averages (vertical lines). Under Moral, we find small insignificant differences between the distributions of donations across complexity levels (Mann-Whitney-U test (MWU); $p = 0.568$). Effects of complexity emerge when comparing donations under Waste. There, we find that distributions of donations under both levels of complexity are different (MWU; $p = 0.076$). These non-parametric results identify general differences between distributions of outcomes across treatments but say little about measures of central tendency (e.g. conditional means) across distributions.

Table 4.1 presents regression analysis of compliance at the intensive margin. We consider

[6] The effects are similar when looking at the number of evaders (χ^2; $p_{morality} < 0.001$; $p_{complexity} = 0.068$).

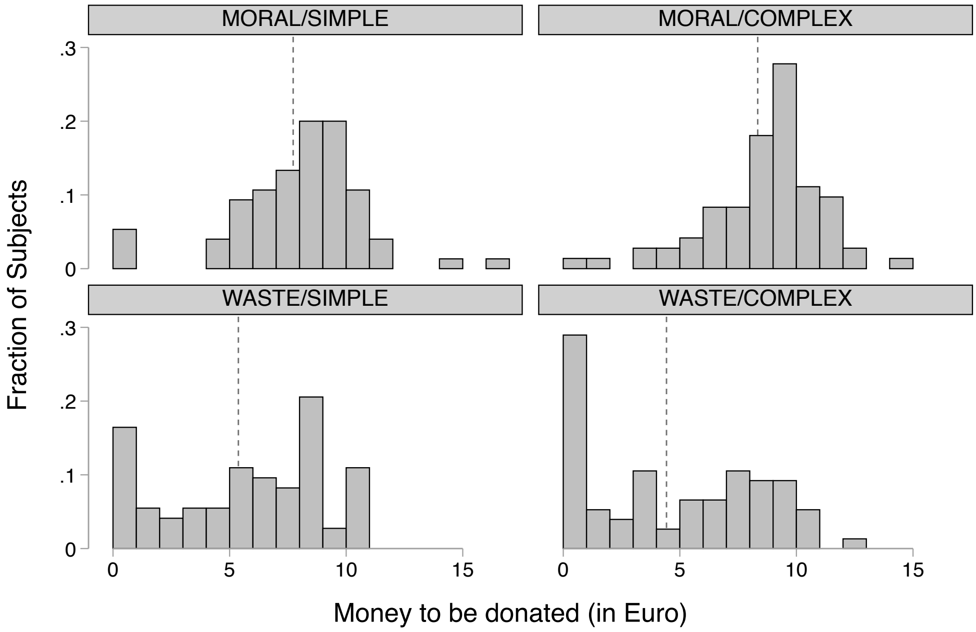

Figure 4.3 Compliance Behavior on the Intensive Margin

Notes: Distributions of money donated to designated charities by treatment. Vertical lines plot average contributions for each treatment.

two related models. The first model, estimated as OLS in column (1) and as Tobit in column (3), regresses donations (in €) on treatment variables and generated income, taking into account or not censoring of donations at 0, respectively.[7] The second model, estimated as OLS in column (2) and as Tobit in column (4), regresses the share of generated income donated on the treatment variables alone.

OLS results suggest that subjects on average donate €2.39 less under Waste/Simple relative to Moral/Simple, a decrease in donations of about 30%. While complexity has a small and insignificant effect under Moral (0.36€), the significant interaction of complexity with Waste suggests that complexity reduced donations by €1.32 (or -18.5%) only when the designated charitable organization was less deserving. This finding is robust to using shares of generated income as the dependent variable (column (2)) or to controlling for censoring of the dependent variable.

These findings are consistent with results on the extensive margin. All suggest that subjects generally donate more to plausibly more deserving charitable organizations, a pure morality effect. The effects of complexity on donations to a deserving charitable organizations are minimal – subjects appear to be willing and able to perform complex calculations

[7] Donations are censored from below at 0 while the share of generated income donated is censored from below at 0 and from above at 1, respectively. One should interpret the p-values on the interaction effect in non-linear models with caution as described by Ai and Norton (2003) and Greene (2010).

	(1) Donation (€)	(2) Donation (%)	(3) Donation (€)	(4) Donation (%)
Waste	-2.394***	-0.133***	-2.575***	-0.145***
	(0.215)	(0.0169)	(0.231)	(0.0188)
Complex	0.360	-0.00142	0.407*	0.00144
	(0.229)	(0.0134)	(0.241)	(0.0141)
Waste × Complex	-1.321***	-0.0549**	-1.556***	-0.0673**
	(0.459)	(0.0261)	(0.539)	(0.0304)
Income	0.189***		0.189**	
	(0.0658)		(0.0746)	
Constant	4.606***	0.468***	4.534***	0.465***
	(1.078)	(0.00944)	(1.224)	(0.0104)
Sigma				
Constant			3.451***	0.206***
			(0.234)	(0.0176)
Observations	296	296	296	296
Model	OLS	OLS	Tobit	Tobit

Table 4.1 Regressions of Compliance Behavior on the Intensive Margin

Notes: Dependent variable in columns (1) and (2) is the amount and share of generated income donated to the designated charity. Here, estimation results do not take into account the censoring of the data. In the last two columns, censoring is taken into account. Clustered standard errors (session level) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

in order to comply with morally justified donation requests. Notably, complexity reduces donations only when the latter are made to less deserving charitable organizations.

4.3.3 Mistakes

261 subjects out of 296 (88%) correctly indicated on their forms that they should donate 50% of their generated income. Figure 4.4 breaks down by treatment the proportion of subjects incorrectly reporting the share of generated earnings they should donate. We find that this proportion is 7% under Moral/Simple and increases to 18% under Waste/Complex. Inaccuracies appear to be higher under Waste than Moral for both levels of form complexity.

Accurate reporting on forms does not automatically lead to compliance, as subjects who correctly indicated on their forms they should keep half of their generated income could do otherwise and leave the experiment with a different share. Selfish subjects for example may decide to keep more of their income than what is prescribed. Selfishness is directly revealed by subjects themselves in this case given they leave behind a clear proof of their understanding of the rules.

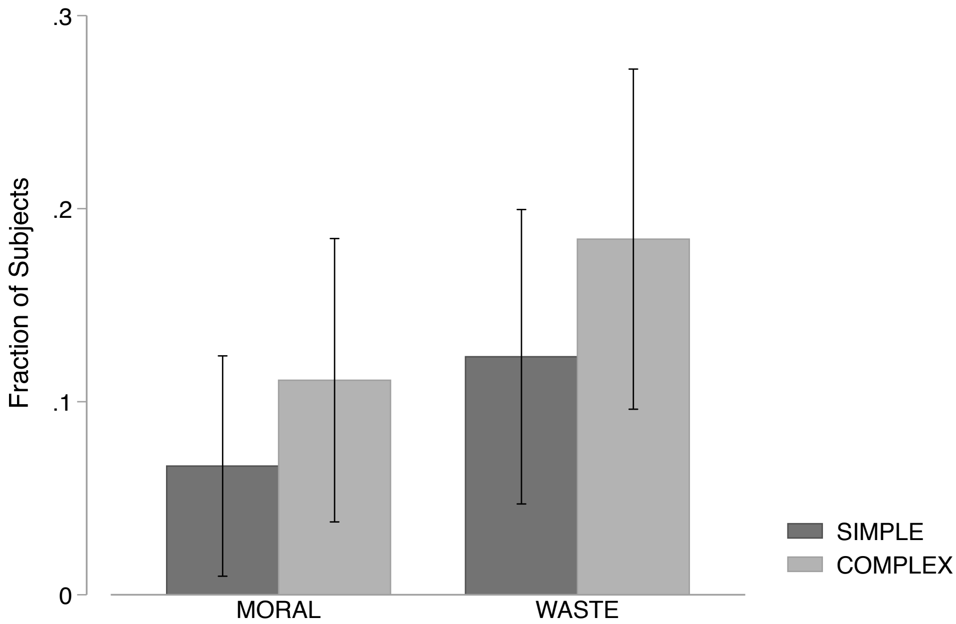

Figure 4.4 Inaccurate Reporting

Notes: Bars show the proportion of subjects inaccurately reporting the share of generated earnings to be donated by treatments. Vertical caps show the 95%-confidence interval that is calculated based on a standard normal distribution.

Figure 4.5 presents by treatment the average share of generated income donated for subjects having inaccuracies on their forms. Overall, we observe that selfishness occurs along with inaccurate reporting; subjects with inaccurate reports donated 14%-points less than accurately reporting subjects (MWU, $p < 0.001$). This pattern appears across all treatments. Admittedly, conditional on forms are complex, we find that inaccurate reporting is not associated with sizeable selfishness when donations are sent to a moral charitable organization – average donations hover near the targeted 50% and are comparable to average donations for subjects having accurately filled out the forms. This is suggestive evidence for the idea that form complexity per se did not lead to systematic deviations of donations from the prescribed rule. A different reporting pattern emerges under Waste/Complex where average donations of subjects with reporting inaccuracies are significantly lower relative to subjects without reporting inaccuracies. Moreover, conditional on reporting inaccuracies, we find that subjects donate significantly less under Waste/Complex relative to Moral/Complex (MWU, $p = 0.002$).

An important question that arises from these observations is whether subjects willingly report self-serving inaccuracies to facilitate non-compliance. Figure 4.6 plots the shares to be donated as well as the shares kept by subjects as they have been reported on the forms. We see no evidence of self-serving mistakes from this figure as one would expect

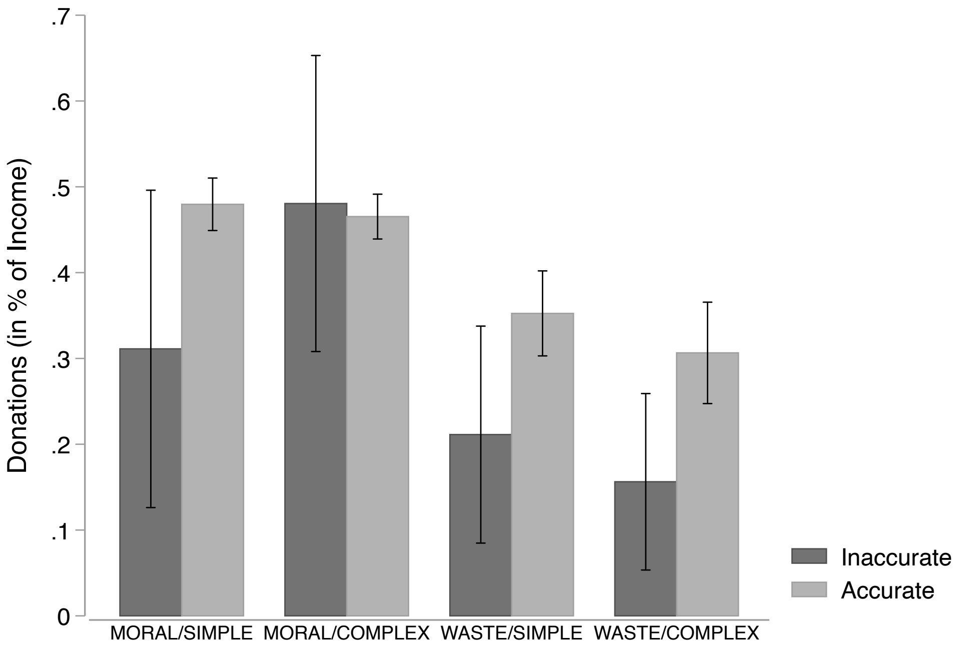

Figure 4.5 Inaccurate Reporting and Donations

Notes: Bars show donations in % of income conditional on inaccurate and accurate reports, respectively. Vertical caps show the 95%-confidence interval that is calculated based on a standard normal distribution.

the reports to be biased favoring the share kept by the subject (or inversely discriminating against donating money to one of the charities). However, it should be noted that strategically creating self-serving mistakes in our setting is costly as it would require, essentially, the subject to work backward through our forms to slip in a convenient mistake at some point. Hence, our reading of the data is that subjects in all treatments managed to figure out how much to donate given their private preferences. If calculating the prescribed compliance level did not work out smoothly and left subjects with uncertainty about the rule, they were able to self-servingly interpret the uncertainty.

4.3.4 Decision Time

We examine decision times to further substantiate the behavioral mechanisms underlying our results. In our design, subjects read the written instructions together before being allowed to open envelopes that included their forms and the instructions related to calculating the shares and the donation procedures. Decision times were measured from that point onwards. Subjects were shown the results from their slider tasks in order to allow them to fill out the forms correctly. We stop the time measurement when the envelopes have been sealed and the subjects end this part of the experiment by clicking on

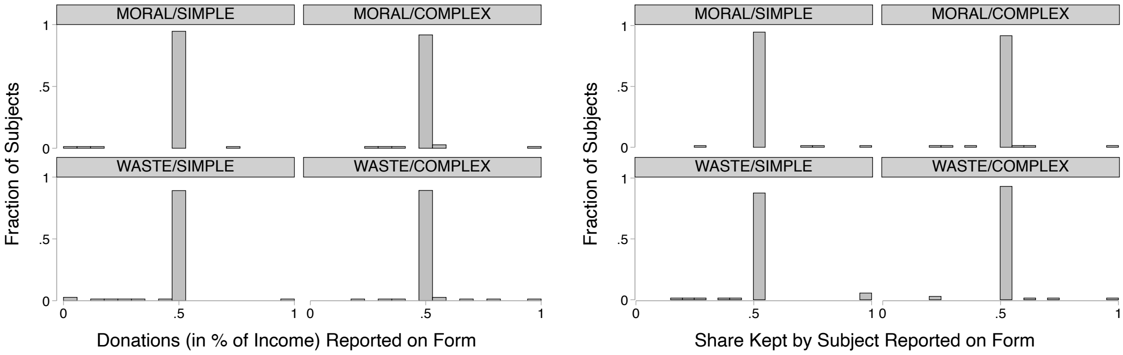

Figure 4.6 Reporting on Forms

Notes: This figure shows histograms of reported donations (in % of income) (left) and the share kept by subjects (right) for all treatments. Subjects that reported shares higher than 100% are not reflected in the histograms. This includes three subjects in the left and four subjects in the right figure.

the respective button. Figure 4.7 presents the distributions of decision times in all four treatments, where horizontal axes measure decision times in seconds.

We observe that decision time distributions under Complex treatments are clearly shifted to the right relative to corresponding distributions under Simple. Table 4.2 presents formal OLS regression analyses of decision times on the main treatment variables of the experiment. We find significantly longer decisions times due to complexity (average increase of 279 seconds) and when the BYC is the recipient (average increase of 32 seconds). There also exists a positive interaction of Waste and Complex of close to 54 seconds, which is significant at the 10% level. Once we control for inaccurate reporting (see column (2)), the treatment interaction gets more precisely measured and remains robust at about 52 seconds. At the same time, not reporting accurately is associated with a time increase of 141 seconds. As shown in column (3), albeit not being statistically significant we observe that decisions are quicker in Waste/Complex when inaccuracies occur as compared to when reporting is accurate.

These observations are in line with our interpretation that filing mistakes induced by complexity are an important facilitator of non-compliance. The decision time data however remains unclear whether the time premium in Waste/Complex is due to strategic filing efforts or the time required to cope with the uncertainty about the compliance rule.

4.3.5 Discussion

Bénabou et al. (2018) develop a model where compliance decisions weight intrinsic duty to

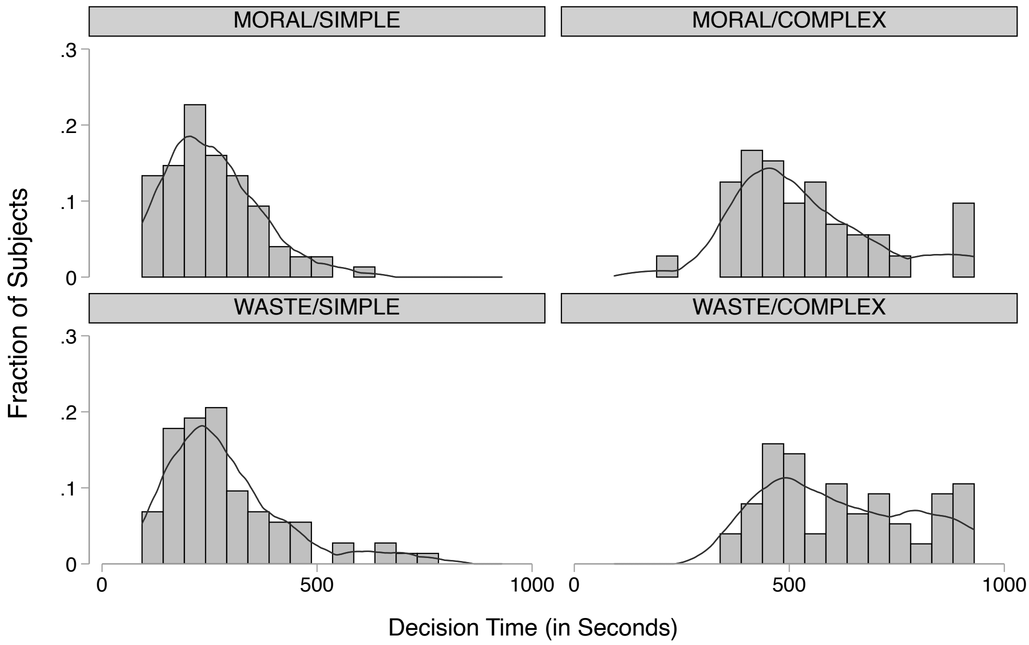

Figure 4.7 Decision Times

Notes: This figure shows histograms of the distributions of decision time in all treatments. The distribution is approximated using a kernel density function.

comply against the moral costs of deviating from the compliance rule.[8] They show that decision-makers can alter the informativeness of the signal that an action sends about their prosocial type, and will do so only if it is effective in maintaining a self- or social-image. Somewhat in line with the model mechanism, compliance in our experiment is significantly related to moral costs of non-compliance. One explanation for our findings is that subjects justify non-compliance in settings where moral costs are low. This conclusion is suggested by the fact that calculation mistakes under the same level of complexity do not bias donations in a specific direction when donations are made to a more deserving charitable organization. However, our subjects did not leave traces behind that could reflect self-serving manipulations of forms as predicted by the model.

In related experiments, Konow (2000) provides evidence on the malleability of fairness perceptions from a set of simple dictator game decisions without relating to complexity as a potential modulating factor. Exley and Kessler (2019) observe that subjects donate less money to a charity when the transferred amount is calculated by, for example, 55+55+55+0 rather than by 55+55+55. When subjects were asked about the result of this sum, they act as if they did not understand how to add a 0 to a sum. However, they

[8] Appendix D.6 presents a simple theoretical framework that illustrates how our design enables the identification of interaction effects. The model however remains agnostic about the underlying mechanism driving the interaction process.

	(1) Time (in Seconds)	(2) Time (in Seconds)	(3) Time (in Seconds)
Waste	31.84*	23.85	12.30
	(17.24)	(14.61)	(19.03)
Complex	279.4***	273.16***	267.3***
	(20.25)	(15.94)	(21.54)
Waste × Complex	53.99*	51.66**	67.09**
	(28.30)	(24.60)	(30.09)
Inaccurate		141.15***	126.1***
		(32.33)	(35.96)
Waste$= 0 \times$ Inaccurate			-126.5
			(96.93)
Complex$= 0 \times$ Inaccurate			-81.30
			(115.6)
(Waste × Complex)$= 0 \times$ Inaccurate			151.5
			(137.6)
Constant	257.3***	247.90***	252.7***
	(13.93)	(13.18)	(16.52)
Observations	296	296	296
Model	OLS	OLS	OLS
R^2	0.535	0.581	0.585

Table 4.2 Regressions of Decision Times

Notes: Dependent variable captures decision time of subjects in seconds. Clustered standard errors (session level) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

have no problem in doing so when the money is split between two charities, i.e. when the tradeoff between money for themselves and money for a charity is eliminated. In contrast to our setting in which the induced complexity was really complex, in Exley and Kessler (2019) it was easier to self-servingly generate mistakes. Exley (2020) uses a related design in which charity performance metrics (in particular the program expense rate) are used by lab participants to construct excuses not to donate. A low program expense rate can be interpreted as the less purposeful appropriation that is related to the yacht club versus the cancer charity. Haisley and Weber (2010) show evidence that experimental subjects in simple ambiguous dictator games have self-serving beliefs about ambiguity which permits justifications to realize unfair allocations in the game. In our experiment, the strong but unbiased correlation between reporting inaccuracies and donations might reflect a similar justification effect that is based on the uncertainty that evolves from complex forms. While selfish behavior in our experiment has been facilitated by filing mistakes, our subjects did not exhibit a tendency to make more self-serving rather than self-hurting mistakes as reported by Leib et al. (2019).

In order to provide more real world context for our lab findings, we exploit data from a representative survey of 1,501 citizens living in the US provided by the PEW Institute. Data contain opinions about tax complexity as well as attitudes towards the fair income share of federal taxes to be paid. We find that 29% of respondents indicating not being bothered by tax complexity report that they pay more than a fair share of their income for taxes. In contrast, more than 50% of respondents indicating being bothered significantly by tax complexity perceive their share of taxes as unfair. This difference remains statistically significant below the 1%-level when controlling for core socio-economic variables (gender, age, income), party preferences, and ideological views.[9] Additional survey evidence from Gallup suggests that about 50% of tax payers in the US perceive their tax payments as wasted money rather than money spent for the public good.[10] This combined evidence suggests that tax complexity may affect attitudes towards tax perception and redistribution, consistent with our experimental findings.

In contrast, charity specific compliance motivations are unlikely to explain our results as we focus on subjects that fill in all forms completely. If the subjects motivation was affected by the appropriation of the taxed money, we would have expected that we see strong responses on the filing margin. We also find no difference across treatments with respect to selective or incomplete fillings (see Appendix D.5). Moreover, it is unlikely that our results are driven by depletion effects that have been shown to increase unethical behavior by reducing self-regulatory resources of experimental subjects (e.g., Mead et al., 2009; Gino et al., 2011). While subjects may have been depleted by filing complex forms, depletion along the moral treatment dimension is unlikely. In addition, longer decision times due to complexity are inconsistent with impulsive decisions that are usually observed in the depletion literature.

4.4 Conclusion

Tax systems serve to achieve a myriad of social and political goals.[11] Achieving these goals simultaneously often requires complex tax codes and filing procedures for many individuals. In return, complexity imposes costs that should be taken into account to determine effectiveness of tax systems. Costs of complex taxation have mostly been as-

[9]Estimation results are available upon request. Similar tax complexity interactions can be observed for the attitudes towards increases or decreases of the federal budget to assist low income individuals in the US and in the world, respectively.

[10]See https://news.gallup.com/poll/232361/less-half-say-taxes-high.aspx.

[11]See Hettich and Winer (1988) for a model how the intricacies of observed tax systems can be viewed as the outcome of optimizing political and economic behavior in the context of, potentially divergent, goals. Hence, tax structure is a system of related parts in equilibrium, not merely a collection of separate and ill-designed components.

sociated with compliance costs (e.g., Benzarti, 2020). We showed that increasing moral costs of non-compliance increases compliance rates. We further documented a significant interaction of morality and complexity effects. In particular, complexity has negative effects on compliance behavior only when moral costs of non-compliance are low. Our data appear consistent with subjects using complexity as a means to decrease psychological costs from non-compliance. Complexity can thus be used to motivate non-compliance and erode the effectiveness of policies which are not perceived as effective or morally justified. This calls for intensified efforts in reducing the overwhelming complexity of tax filing processes in many countries.

Appendix A

Cooperation in a Company: A Large-Scale Experiment

A.1 Variable Overview

Variable	Scale	Description	Details
team membership	nominal	Unique team identifier (from ORG structure)	
team size	ratio	Number of team members	
age	ratio	Age of employee	
gender	nominal	Gender of employee	
seniority	ratio	Seniority of employee (in years)	
job function	nominal	Twelve functional areas (departements) which consists of clusters of several job families based on generic job content	Communications, Development, Education and Training, Finance, Administration, Human Resources, Information Technology, Marketing, Sales, Consulting, Not assigned
career	ordinal	Nine career level of employee (describes contribution based upon business results, accountability, complexity, experience and communication)	Not specified for reasons of discretion
leader	binary	Leadership responsibility	Yes/No
pay scheme	nominal	Employees pay scheme	Either company performance pay or individual performance pay
wage	ratio	Yearly wage before taxes	
financial awards	ratio	Amount of money received in a year	
received recognition	ratio	Number of peer-to-peer awards received for being cooperative (received)	
sent recognition	ratio	Number of peer-to-peer awards received for being cooperative (sent)	

Table A.1 Variables Collected from Company Records

Variable	Scale	Description	Details
contribute	ratio	Unconditional contribution	
x-contribute	ratio	Contribution conditional on x contributed by other team members	
belief contribute	ratio	Belief about average contribution of the other team members	
with mpcr variation	ratio	Only in 2019: *contribute* and *belief contribute* for mpcr= 0.3 and mpcr= 1.2	Strategy method, within subject
donation	nominal	Only in 2017: earnings from experiments transferred to individual account or to a charity	Personal bank account, Deutsche Aidshilfe, rzte ohne Grenzen, World Wide Fund For Nature (WWF), SOS Kinderdorf, Amnesty International

Table A.2 Variables Collected from the Experiments

Variable	Scale	Description	Details
team cooperation	ordinal	Need for cooperation among team members	
team cohesion	cardinal	Perception of team cohesion	
team stability	cardinal	Perception of staff stability within the team	
neg. competiveness	ordinal	Perception of negative competitive pressure among team members	
pos. competiveness	ordinal	Perception of positive competitive pressure among team members	
stress	cardinal	Perceived chronic stress	Individual average score
big five	cardinal	big five personality measure	Individual average score (for each personality trait)
neg. reciprocity	ordinal	Social preference measure indicating the participants tendency for negative reciprocity	
pos. reciprocity	ordinal	Social preference measure indicating the participants tendency for positive reciprocity	
trust	ordinal	Social preference measure indicating the participants trust	
competitive attitude	cardinal	The participants individual competitive attitude	Individual average score
children	binary	Indicating whether the participant has children or not	
friends	ratio	The participants number of friends	
complement	nominal	In which business model is the employee working? Cloud model requires much more cooperation than customer model	Cloud, Customer, Neither

Table A.3 Variables Collected from the Survey

A.2 Instructions

You are a member of a group of three, consisting of anonymous participants in this study. All participants are randomly selected employees of [COMPANY]. The combination into groups of 3 occurs randomly. The payouts for you and the other group members in this section depend on your decisions and the decisions of the other members of your group.

Decision-making situation
Each member of the group must decide on the use of 10 tokens each. You and the other group members can put the 10 tokens into a private account, or you can deposit them in whole or in part into a common account. Any tokens that you do not deposit into the common account are automatically added to your private account.

Income from the private account
You earn exactly one euro for each token you put in your private account. For example, if you put 4 tokens into your private account, you will earn exactly €4 from your private account. No one but you receives income from your private account.

Income from the common account
For each token that is added to the common account, you will receive €0.5. The other two group members also each receive €0.5 for each token you contribute. Conversely, you also earn money from the contributions of the other two group members to the common account. The income of each member from the common account is determined as follows:

Individual income from the common account = Sum of the contributions of all three group members to the common account times 0.5

For example, if the sum of all three group members' contributions to the common account results in 30 tokens, then you and the other two group members each receive 30 x 0.5 = €15 from the common account. If the three group members pay a total of 10 tokens into the common account, you and the other two group members receive 10 x 0.5 = €5 each from the common account.

Total income Your total income is the sum of your income from your private account and your income from the common account. So:

Income from the private account (= 10 - contribution to the common account) + income from the common account (=0.5 x sum of contributions to the common account) = Total income

As described above, you can use 10 tokens to fund your private account and the common account. Each group member has to make two types of contribution decisions, which we will refer to below as the contribution and the contribution table. You can find a detailed description of your entries on the entry screens.

A.2.1 Comprehension Questions

Please answer the following questions to ensure that you have understood the instructions of the experiment. If you are unsure, you can return to the instructions by clicking on "Back".

1. Assume that none of the group members (even you yourself) pay a contribution into the group account.
   - How high is your total income?
   - How high is the respective total income of the other two group members?
2. Assume that all three group members (also you yourself) each pay a contribution of 10 tokens into the group account.
   - How high is your total income?
   - How high is the respective total income of the other two group members?
3. Assume that you deposit 0 tokens into the common account and that the other two members of your group deposit 10 tokens each.
   - How high is your total income?
   - How high is the respective total income of the other two group members?
4. Assume that you pay 10 tokens into the common account and the other two members of your group each pay 0 tokens.
   - How high is your total income?
   - How high is the respective total income of the other two group members?

A.2.2 Contribution Decisions

When choosing the contribution to the common account, you determine how many of the 10 tokens you want to deposit into the common account. The deposit to your private account is automatically the difference between 10 tokens and your contribution to the common account.

- Please enter the amount you would like to pay into the common account (any whole-number value between and including 0 and 10 is possible): ...

Now you will be asked to fill in a contribution table. In the contribution table, you should specify how many tokens you want to pay into the common account for each possible (rounded) average contribution of the other two group members to the common account. So, depending on how much the others contribute on average, you must define your own contribution decision. For each average contribution of the other two group members, please indicate the amount you would like to pay into the common account (any whole-number value between and including 0 and 10 is possible; of course, you can also enter the same amount several times):

What is your contribution to the common account if...

- ... the other two group members deposit an average of 0 tokens.
- ... the other two group members deposit an average of 1 tokens.
- ... the other two group members deposit an average of 2 tokens.
- ... the other two group members deposit an average of 3 tokens.
- ... the other two group members deposit an average of 4 tokens.
- ... the other two group members deposit an average of 5 tokens.
- ... the other two group members deposit an average of 6 tokens.
- ... the other two group members deposit an average of 7 tokens.
- ... the other two group members deposit an average of 8 tokens.
- ... the other two group members deposit an average of 9 tokens.
- ... the other two group members deposit an average of 10 tokens.

Help option: The numbers in the left column are the possible (rounded) average contributions of the other two group members to the common account. You now have to specify how many tokens you want to deposit into the common account for each slider, provided that the others contribute the specified amount on average. You have to make an entry in each field. For example, you are to specify how much you contribute to the common account if the other group members deposit an average of 0 tokens into the common account; how many tokens you contribute if the others contribute an average of 1 token or 2 tokens or 3 tokens, and so on. You can enter any whole-number contribution from 0 tokens to 10 tokens in each field and, of course, the same amount several times.

A.2.3 Incentive Compatibility

Payout relevance of your decisions
After all study participants have made their decisions, one member is randomly selected in each group of 3. For the randomly selected member, only the contribution table filled in by him/her is relevant for decision making and payout. For the other two group members who have not been selected, only the contribution is relevant for decision-making and payout. The average of the two contributions (rounded to the next whole number) then determines the relevant conditional contribution from the third member's contribution table. Of course, you do not yet know which of your contribution decisions will be randomly selected. You must therefore carefully consider both types of contribution decisions, as both can become relevant to you.

The following graphic (Figure B.1) is intended to visualize the decision-making situation. For the randomly selected person on the right, the conditional contribution from the contribution table is relevant. For the other two group members, the contribution is relevant for payout.

A.2.4 Belief Elicitation

In addition to your earnings from your private and common account, you will receive a further payout for estimating the average contribution of the other two members of your group to your common account. Your payout will depend on how accurately you estimate the actual average contribution of your two group members. If you are exactly right, you will receive an additional €5. If your estimate differs by 0.5 or more tokens from the actual average contribution, you will receive €0. Please enter a number from 0

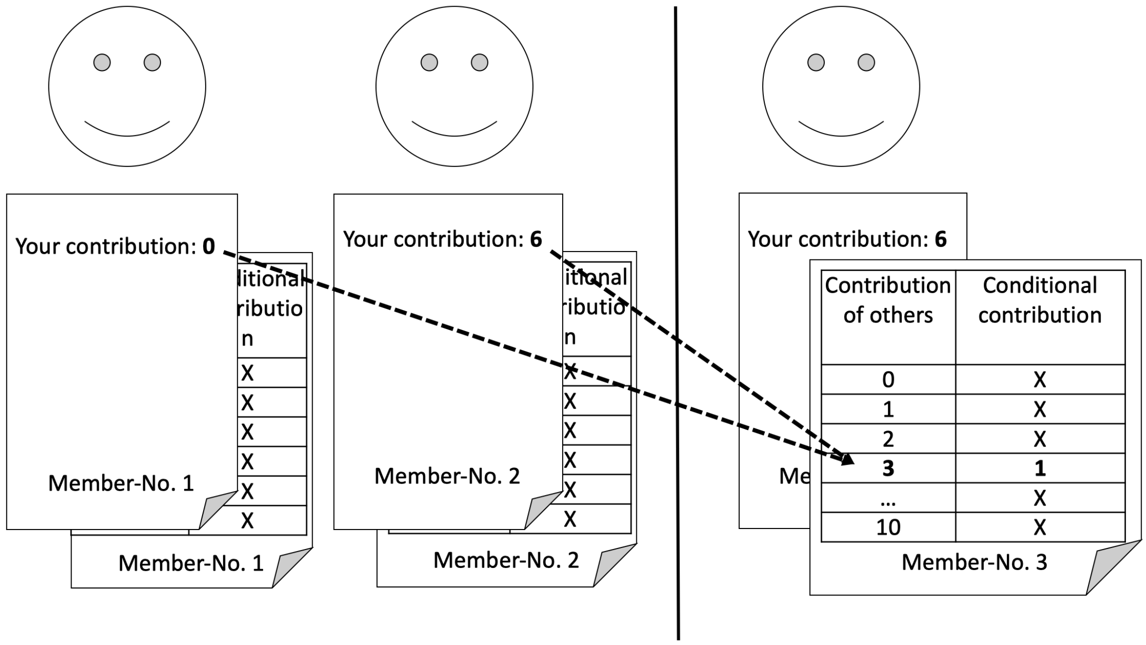

Figure A.1 Incentive Compatibility

to 10 (each number is allowed in steps of 0.5).

What do you think is the average amount of tokens your two group members contribute to the common account?

- ... Average contribution of the other two members of your group

A.3 Communication and Coordination of Employees

Employees could interrupt the experiments and continue at a later point in time. On average, employees finished the experiment and survey within approximately one and a half days (mean=1.35 days). While employees in a public goods game group were anonymously selected and matched, one might be concerned about communication and coordination during the experiment as some teams in the company are seated in shared offices (max four team members per office). To alleviate this concern, we observe no correlation between contribution behavior, beliefs and attitudes of employees with respect to the variance of finishing times within work teams (Spearman Correlations; uncond. contribution, $\rho = -0.004$, $p = 0.905$; belief about others' uncond. contribution, $\rho = -0.006$, $p = 0.853$; mean cond. contribution, $\rho = 0.008$, $p = 0.827$).

A.4 Overview of Public Goods Game Measures

Our employee sample appears to be very cooperative as can be seen from Table 1.2. In the unconditional contribution decision, they contribute on average 7.9 Tokens (which corresponds to 79% of the endowment) in the public good. The average belief about the public good contribution of the other group members equals 6.7 Tokens. The difference in actual contributions and beliefs is statistically significant at the 1%-level (Wilcoxon Sign Rank Test, $p < 0.001$). Reassuringly, we observe very similar responses in the public goods games when comparing the variables collected from the experiments in 2017 and 2019. This holds for the data presented in Table A.4 – see column "Comparison".

	All			Experiment 2017			Experiment 2019			Comparison
	Count	Mean	SD	Count	Mean	SD	Count	Mean	SD	P-Value
Unconditional Contributions	910	7.89	2.93	438	7.90	2.97	472	7.88	2.90	0.890
Belief about Others' Contributions	910	6.70	2.78	438	6.73	2.82	472	6.67	2.75	0.751
Mean Conditional Contribution	910	5.30	2.25	438	5.26	2.19	472	5.33	2.29	0.607
Slope Parameter	910	0.71	0.43	438	0.71	0.43	472	0.70	0.43	0.818

Table A.4 Overview of Public Goods Game Variables by Wave

Notes: P-values rely on two-sample Mann-Whitney-U tests.

We observe that cooperative attitudes are highly predictive for unconditional contributions, also when we control for beliefs about other's contributions (see Table A.5). Net-Givers contribute more than Matchers and Matchers contribute more than Net-Takers. Both differences are highly statistically significant.

The scatter plot in Figure A.2 shows a significant variation in the average conditional contributions and the reciprocity parameter of employees. The size of the bubbles represents the frequency of the observed combination of mean conditional contribution and reciprocity. There are several mass points that stand out.

Next to our cooperative attitudes, we also classify cooperation types as described by Fischbacher et al. (2001) and Fischbacher and Gächter (2010). These types are also visible in the scatter plot. First, we observe employees that behave like perfect conditional cooperators ($[1, 5]$). Secondly, there are clusters of employees whose contributions are independent of the contribution schedule. They are either contributing nothing (free-riders) or they contribute a strictly positive amount (unconditional cooperators). Most of the unconditional cooperators contribute all their endowment. Thirdly, imperfectly condi-

	Uncond. Contribution
Belief About Others' Undcond. Contribution	0.642***
	(0.0269)
Net-Taker	0
	(.)
Matcher	1.196***
	(0.196)
Net-Giver	1.536***
	(0.198)
Dummy (2019)	0.00693
	(0.137)
Constant	2.520***
	(0.282)
Observations	910
R^2	0.509
Model	OLS

Table A.5 Determinants of Unconditional Contributions

Notes: Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

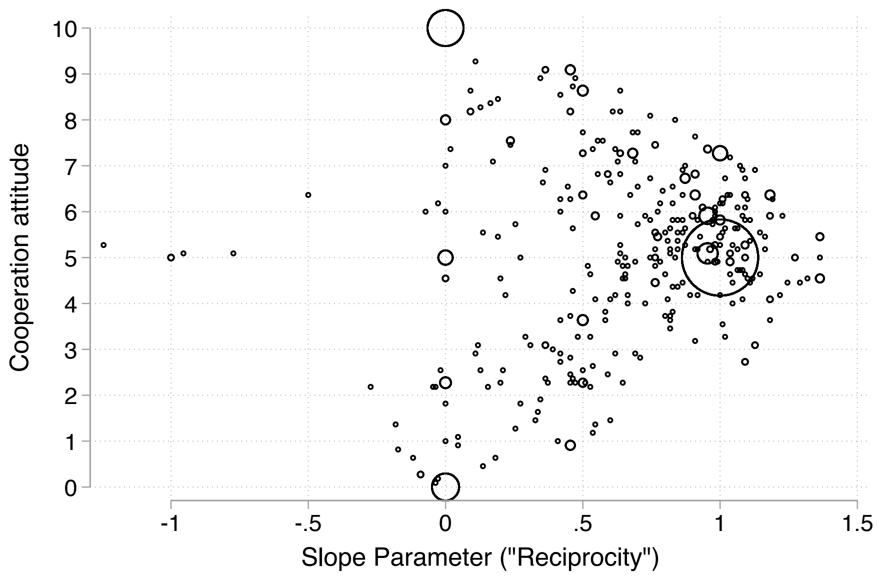

Figure A.2 Relationship Between Mean Cond. Contributions and Slope Parameter

Notes: The graph contains data from all participating employees. Bubble sizes show the frequency of the combination of both variables. Reciprocity is the slope parameter from an OLS regression between an employees' conditional contributions and the contribution schedule.

tional cooperators are split in two groups, conditional cooperators with a self-serving bias (mean unconditional contribution below 5) and conditional cooperators with an other-serving bias (mean unconditional contributions above 5). The remaining employees are classified as Others.

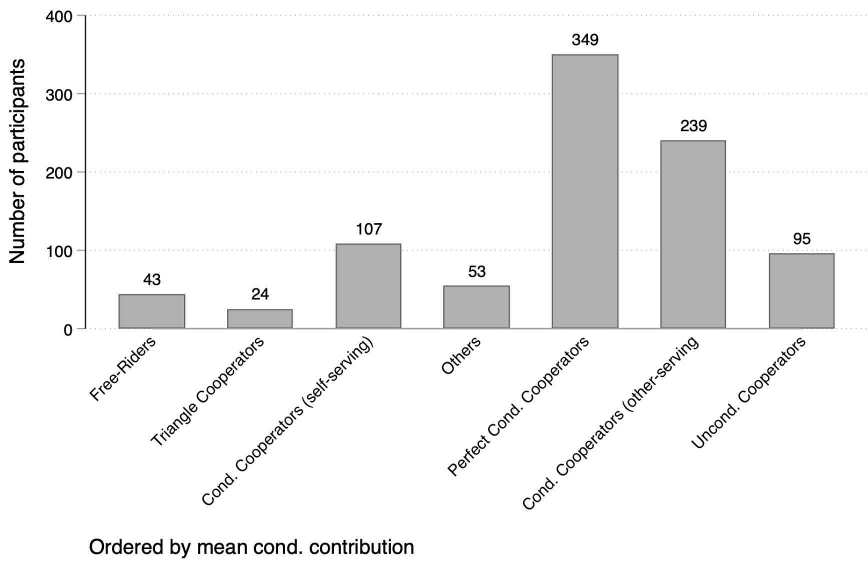

Figure A.3 Distribution of Cooperation Types

Notes: Bars show the fraction of all participating employees that belong to a particular cooperation type. Bars are ordered by mean conditional contributions.

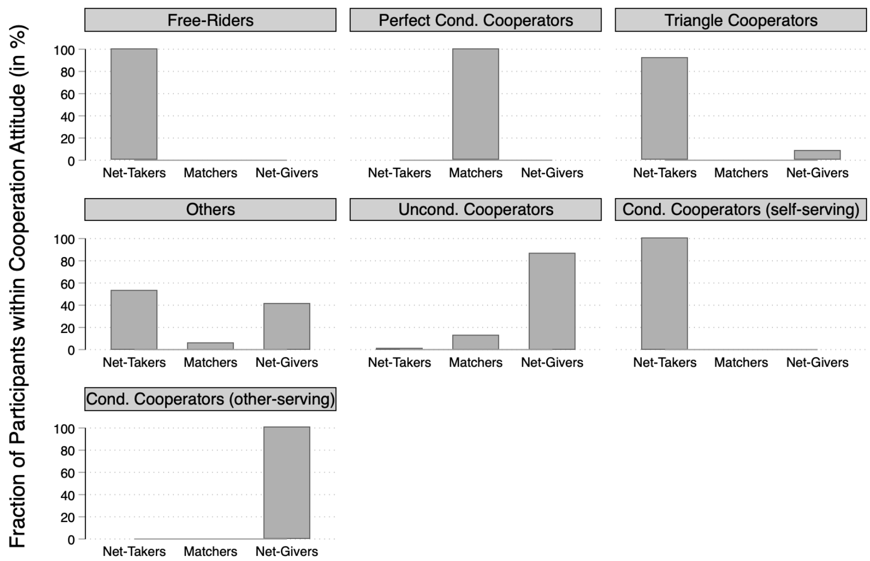

Figure A.4 Cooperation Types and Cooperative Attitudes

Notes: Bars show the fraction of participating employees that belong to a particular cooperative attitude.

Figure A.3 shows an overview of all types and Figure A.4 relates our cooperative attitudes and the cooperation types. Cooperative attitudes subsume the classification types reasonably well. We use cooperative attitudes because they prove handier for the statistical analysis.

Table A.6 shows an inverse U-shape relationship between cooperative attitudes and reciprocity.

	Slope Parameter
Net-Taker	0
	(.)
Matcher	0.497***
	(0.0328)
Net-Giver	0.134***
	(0.0331)
Constant	0.456***
	(0.0263)
N	910
R^2	0.238
Model	OLS

Table A.6 Reciprocity and Cooperative Attitudes

Notes: Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.5 Correlates of Cooperative Attitudes

In Table A.7, we present correlations of cooperative attitudes with socio-demographic characteristics and behavioral measures. We account for the categorical scale of the dependent variable by using a Multinomial Logit Regression Model.

	Cooperative Attitude		
	Matcher	Net-Giver	Net-Taker
Age	0.00196	0.0252**	*Base Category*
	(0.0108)	(0.0110)	
Female	-0.166	-0.418**	
	(0.192)	(0.200)	
High Education	0.102	-0.0451	
	(0.201)	(0.205)	
Patience	0.0191	0.0226*	
	(0.0131)	(0.0136)	
Competitiveness	-0.286***	-0.391***	
	(0.103)	(0.106)	
Distrust	-0.175*	-0.278***	
	(0.103)	(0.108)	
Positive Reciprocity	0.368**	0.442***	
	(0.153)	(0.161)	
Negative Reciprocity	0.290***	0.0965	
	(0.103)	(0.109)	
Dummy (2019)	-0.00368	-0.198	
	(0.198)	(0.201)	
Constant	-1.431	-1.716	
	(1.066)	(1.115)	
Observations	905		
Pseudo R^2	0.032		
Model	MnLogit		

Table A.7 Cooperative Attitudes, Socio-Demographics and Behavioral Measures

Notes: Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; Five subjects missing because they did not insert information on their socio-demographic status. High Education is an indicator for higher than median education subsuming two out of five education categories.

First, we observe an indication for age being positively related to the cooperativeness of employees. Older employees are significantly more likely to be Net-Givers rather than Net-Takers. The share of Matchers is relatively stable across age cohorts. Second, female employees are less frequently Net-Givers than Net-Takers and, again, the share of Matchers is very similar. Marginal effect calculations show that female employees are about 7%-points more likely to be Net-Takers rather than Net-Givers than male employees are. Third, the competitiveness index correlates with cooperative attitudes. Intuitively, employees are more likely to be Net-Takers the more competitive they are. Moreover, we find that the agreement to the statement “You can’t trust strangers anymore” is highly predictive for the cooperative attitude. The likelihood of being a Net-Taker decreases with reported distrust in strangers. Finally, we observe positive correlations between survey measures for positive and negative reciprocity (agreement with "When someone does me a favor, I am willing to return it" and "If I am treated unjustly, I will take revenge at the first occasion, even if there is a cost to do so", respectively) and cooperative attitudes – again, in the expected positive direction.[1]

In Table A.8, we present the correlations between cooperative attitudes and structural variables from the company context. Here, the main observation is that the cooperativeness of employees is less pronounced in the individual performance pay scheme. While we classify 20% of participants in the company performance pay scheme as Net-Taker, the respective share increases to 27% in the individual performance pay. This increase in the share of Net-Takers comes along with a decrease in the share of Matchers (from 41% to 35%). The share of Net-Givers is not significantly different between incentive schemes. We do not observe significant differences in the distribution of cooperative attitudes with respect to career levels, leadership responsibility, seniority, or the team work production function.

Lastly, we use a short form of the big five personality trait questionnaire validated by Rammstedt et al. (2013) from our online survey. The traits consist of extraversion, agreeableness, conscientiousness, neuroticism, and openness. Table A.9 shows the correlation between our cooperative attitude classification and the five traits. Net-Takers are significantly more extroverted and neurotic than Net-Givers, and more conscientious than Matchers.

[1] Other studies report that female subjects (both employees and students) are more cooperative (e.g., Charness and Villeval, 2009). Low cooperativeness of women compared to men in our context could be related to the selection of women working in a male-dominated work environment.

	Cooperative Attitude		
	Matchers	Net-Givers	Net-Takers
Female	-0.210	-0.415**	*Base Category*
	(0.200)	(0.206)	
High Education	0.127	0.0114	
	(0.209)	(0.210)	
Seniority	0.00382	0.0182	
	(0.0160)	(0.0161)	
Low Career Level	0.152	-0.105	
	(0.433)	(0.459)	
Medium Career Level	0.0924	0.403	
	(0.333)	(0.343)	
High Career Level	0	0	
	(.)	(.)	
Leader	-0.133	-0.0317	
	(0.322)	(0.319)	
Ind. Performance Pay	-0.511**	-0.370	
	(0.232)	(0.231)	
Cloud	0	0	
	(.)	(.)	
Customer	-0.163	-0.139	
	(0.218)	(0.222)	
Neither	-0.387*	-0.0336	
	(0.234)	(0.231)	
Dummy (2019)	-0.0606	-0.303	
	(0.218)	(0.217)	
Constant	0.792	0.714	
	(0.539)	(0.548)	
Observations	861		
Pseudo R^2	0.014		
Model	MnLogit		

Table A.8 Cooperative Attitudes, Socio-Demographics and Company Variables

Notes: Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; 49 subjects are missing be-cause they did not insert information of their socio-demographic status or there was no wage data available. High Education is an indicator for higher than median education subsuming two out of five education categories. Career levels subsume several categories in each presented category.

	Cooperative Attitudes		
	Matchers	Net-Givers	Net-Takers
Age	0.00814	0.0316***	*Base Category*
	(0.0106)	(0.0109)	
Female	-0.0421	-0.291	
	(0.192)	(0.200)	
High Education	0.0749	-0.145	
	(0.201)	(0.204)	
Extraversion	-0.136	-0.203**	
	(0.0927)	(0.0949)	
Agreeableness	0.0346	0.0957	
	(0.122)	(0.125)	
Conscientiousness	-0.272**	-0.129	
	(0.138)	(0.143)	
Neuroticism	-0.134	-0.292***	
	(0.103)	(0.107)	
Openness	0.0323	0.149	
	(0.0984)	(0.101)	
Dummy (2019)	-0.0274	-0.258	
	(0.195)	(0.197)	
Constant	1.972**	0.803	
	(0.984)	(1.021)	
Observations	906		
Pseudo R^2	0.018		
Model	MnLogit		

Table A.9 Cooperative Attitudes, Socio-Demographics and Personality Traits

Notes: Standard errors in parentheses; * p < 0.10, ** $p < 0.05$, *** $p < 0.01$; Four subjects missing because they did not insert information on their socio-demographic status. High Education is an indicator for higher than median education subsuming two out of five education categories.

A.6 Description of Outcome Variables

In the following, we provide descriptive analyses of our main outcome variables. Company variables stem from the records as of 12/31/2017 for employees that were invited to participate in the experiments in 2017. For employees invited to the second experiment, we use record data as of 12/31/2018.

Table A.10 shows the data availability for our main outcome variables. We have data on recognition awards from 2017 for employees that participated in 2017 and the data from 2018 for the participants from 2019. Wage data covers 2016/2017 and 2017/2018 for the employees in the different roll-out phases, respectively. This allows us to look at changes in wage over time. We do not have information on financial awards in 2018 for employees from the first experiments due to data restrictions at the company. In addition, the company-wide budget for the financial award allocation differed strongly between 2017 and 2018 such that there is low comparability.

	Wage			**Financial Awards**		**Recognition Awards**	
	2016	2017	2018	2017	2018	2017	2018
Participants from Experiments 2017	✓	✓	X	✓	X	✓	X
Participants from Experiments 2019	X	✓	✓	✓	✓	X	✓

Table A.10 Overview of Record Data Used as Outcome Variables

Notes: Table shows the data availability of our main outcome variables for employees that could participate in 2017 and 2019, respectively. The variables are retrieved from the company records at the 12/31/2017 or the 12/31/2018, respectively.

Wages. Between 2016 and 2018, participating employees received an average yearly wage increase (in percent of the previous year) of 4.5% with a standard deviation of 8.61% within the range of -63.4% and 72.5%. Calculated as a full-time position equivalent, we observe an average increase of 4.2% with a standard deviation of 3.58% within the range of -8.5% and 33.9%. A full-time wage equivalent equals the nominal wage dived by the part-time share. For example, if an employee receives a wage of €50,000 but works part-time on a 50% position, the full-time equivalent is 50,000/50% = €100,000. Here, we assume a linear relation between the part-time parameter and the wage level which might not be true. In our analyses, we rely on the nominal compensation changes and levels (including potential variations in the part-time parameter).

Financial awards. We measure the award value in percent of the wage in 2017. In this year, conferred awards were worth up to 30% of the yearly wage. The average award payment was about 6% (standard deviation of 5.5%). About 60.4% of employees received an award payment larger than 0.

Recognition awards. In total, we observe 354 recognition awards received by 225 (38.90%) employees and 274 awards sent by 102 (11.21%) employees in 2017 and 2018.[2] Conditional on sending at least one award, we observe that employees sent up to 20 awards with the median number being 2 and a mean of 2.69. Conditional on receiving at least one award, employees received up to 7 awards with the median number being 1 and an average of 1.57.

[2]The number of received awards and the number of sent awards do not need to equalize because we only have a subsample of employees and awards can, of course, be sent to non-participating employees.

A.7 Part-Time Variations and Financial Rewards

Table A.11 shows robustness analyses of our main effects with regard to the part-time share of employees.

	(1) Wage Increase (in %)	(2) Fin. Award Payment (in %of Wage)	(3) FTE-Wage Increase (in %)
Net-Takers	0	0	0
	(.)	(.)	(.)
Matchers	0.00568	0.000596	0.00128
	(0.00403)	(0.00506)	(0.00319)
Net-Givers	0.00413	-0.00312	0.00218
	(0.00289)	(0.00547)	(0.00275)
Ind. Perf. Pay	0.000506	-0.00326	-0.0000508
	(0.00599)	(0.00819)	(0.00580)
Net-Takers ×	0.0208***	0.0270***	0.0148**
Ind. Perf. Pay	(0.00773)	(0.0103)	(0.00734)
Matchers ×	0.00898	-0.00361	0.0128
Ind. Perf. Pay	(0.00855)	(0.00928)	(0.00818)
Net-Givers ×	0	0	0
Ind. Perf. Pay	(.)	(.)	(.)
Δ(Part-time Share)	0.0127***		
	(0.000541)		
Part-time Share		0.0000577	
		(0.000215)	
Constant	0.0909***	0.0145	0.0979***
	(0.0126)	(0.0273)	(0.0123)
b[Matchers] -b[Net-Givers]	p=0.655	p=0.412	p=0.737
b[Matchers | IPP] -b[Net-Takers | IPP]	p=0.218	p=0.002	p=0.816
Socio Demographics	✓***	✓***	✓***
Company Controls	✓	✓	✓
Career Dummies	✓***	✓***	✓***
Dep. Dummies	✓***	✓**	✓***
Observations	831	857	831
R^2	0.817	0.244	0.198
Model	OLS	OLS	OLS

Table A.11 Regressions of Financial Appreciation Controlling for Part-Time Effects

Notes: Standard errors clustered on team-level in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; For wage increases, we use data from 2016/2017 for participants in 2017 and data from 2017/2018 for participants in 2019. We use the value of financial award payments received in 2017 in percent of the 2017-wage level. *FTE-Wage Increase* is the full-time equivalent of wage increases. Asterisks for the control variables show the test result from an F-Test, testing the joint difference from zero.

A.8 Wage Levels and Cooperative Attitudes

Columns (1) to (3) of Table A.12 show that we find no significant relationship between wage levels and cooperative attitudes. In columns (4) to (6), we additionally control for an interaction of cooperative attitudes and age. It can be seen that in these regressions, Net-Takers and Matchers earn less than Net-Takers but that this effect decreases with age. This is likely related to the explanations mentioned by us in the main text such as in-/outflux of employees and leveling effects of collective bargaining agreements.

	(1) Ln(Wage) 2016	(2) Ln(Wage) 2017	(3) Ln(Wage) 2018	(4) Ln(Wage) 2016	(5) Ln(Wage) 2017	(6) Ln(Wage) 2018
Net-Takers	0	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)	(.)
Matchers	-0.00136	-0.000917	-0.0217	-0.0360	-0.123	-0.402**
	(0.0288)	(0.0188)	(0.0244)	(0.120)	(0.0950)	(0.177)
Net-Givers	0.0161	0.00706	-0.0123	-0.0167	-0.154*	-0.300*
	(0.0276)	(0.0156)	(0.0218)	(0.112)	(0.0911)	(0.174)
Age	0.0101***	0.0108***	0.0128***	0.00956***	0.00840***	0.00758***
	(0.00197)	(0.00123)	(0.00161)	(0.00271)	(0.00188)	(0.00287)
Net-Takers × Age				0	0	0
				(.)	(.)	(.)
Matchers × Age				0.000810	0.00276	0.00812**
				(0.00278)	(0.00202)	(0.00353)
Net-Givers × Age				0.000758	0.00358*	0.00618*
				(0.00253)	(0.00199)	(0.00355)
Constant	10.49***	10.40***	9.965***	10.51***	10.51***	10.24***
	(0.107)	(0.0788)	(0.213)	(0.130)	(0.0934)	(0.218)
b[Matchers] -b[Net-Givers]	p=0.378	p=0.604	p=0.655	p=0.853	p=0.719	p=0.507
b[Matchers | Age] -b[Net-Givers | Age]	.	.	.	p=0.983	p=0.647	p=0.516
Socio Demographics	✓***	✓***	✓***	✓***	✓***	✓***
Company Controls	✓***	✓***	✓***	✓***	✓***	✓***
Career Dummies	✓***	✓***	✓***	✓***	✓***	✓***
Dep. Dummies	✓***	✓***	✓***	✓***	✓***	✓***
Observations	367	857	467	367	857	467
R^2	0.785	0.752	0.746	0.785	0.752	0.749
Model	OLS	OLS	OLS	OLS	OLS	OLS

Table A.12 Regressions of Wage Levels on Cooperative Attitudes

Notes: Standard errors clustered on team-level in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; For wage increases, we use data from 2016/2017 for participants in 2017 and data from 2017/2018 for participants in 2019. We use the value of financial award payments received in 2017 in percent of the 2017-wage level. Asterisks for the control variables show the test result from an F-Test, testing the joint difference from zero.

A.9 Survey Outcomes and Cooperative Attitudes

In Table A.13, we show OLS regression models of both variables on the share of Net-Takers, Matchers, and Net-Givers in a work team, estimated using regressions with analytical weights to account for team-specific participation rates . We detect no statistically relevant relationship to the perception of team stability as shown in column (1). However, in column (2), we find that members of teams that perceive themselves as being in a more cohesive team tend to be more cooperative in the experiment.

	(1) Team Stability	(2) Team Cohesion
Net-Taker Share	-0.0438	-0.192
	(0.423)	(0.397)
Matcher Share	-0.0709	-0.103
	(0.299)	(0.285)
Net-Giver Share	-0.282	0.651**
	(0.338)	(0.272)
Constant	2.950***	4.333***
	(0.258)	(0.255)
Socio-Demographics	✓	✓**
Company Controls	✓***	✓
Career Dummies	X	X
Department Dummies	X	X
Observations	299	299
R^2	0.076	0.046
Model	WLS	WLS

Table A.13 Regressions of Team Cohesion on Team Composition

Notes: Robust standard errors in parentheses; * $p < 0.10$, $** p < 0.05$, *** $p < 0.01$; We use OLS with analytical weights that emphasize averages of teams that participated with a higher share of team members. We control for gender and age composition as well as average seniority. We do not control for career levels or function compositions because of the large number of different categories.

Appendix B

Cooperation, Free-Riding, and the Signaling Value of Incentives: An Experiment in a Company

B.1 Variable Overview

Variable	Scale	Description	Details
age	ratio	Age of employee	
gender	nominal	Gender of employee	
seniority	ratio	Seniority of employee (in years)	
job function	nominal	Twelve functional areas (departments) which consists of clusters of several job families based on generic job content	Communications, Development, Education and Training, Finance, Administration, Human Resources, Information Technology, Marketing, Sales, Consulting, Not assigned
career	ordinal	Nine career level of employee (describes contribution based upon business results, accountability, complexity, experience and communication)	Not specified for reasons of discretion
pay scheme	nominal	Employees pay scheme	Either company performance pay or individual performance pay

Table B.1 Variables Collected from Company Records

Variable	Scale	Description	Details
Employees			
contribute	ratio	Unconditional contributions with and without the incentive in place	
x-contribute	ratio	Contribution conditional on x contributed by other team members with and without the incentive in place	
belief contribute	ratio	Belief about average contribution of the other team members with and without the incentive in place	
manager choice	ratio	Belief about share of managers that select the incentive	
manager belief	ratio	Belief about managers' expectation about unconditional contributions of employees	
Managers			
incentive choice	binary	Decision about whether to set the incentive	
belief contribute	ratio	Belief about average contribution of employees with and without the incentive in place	
2nd order belief	ratio	Belief about employees' beliefs about contributions of others with and without the incentive in place	

Table B.2 Variables Collected from the Experiment

Variable	Scale	Description	Details
altruism	ordinal	Social preference measure indicating the participants tendency for altruistic behavior	
neg. reciprocity	ordinal	Social preference measure indicating the participants tendency for negative reciprocity	
pos. reciprocity	ordinal	Social preference measure indicating the participants tendency for positive reciprocity	
math	ordinal	Measure of perceived math skills	
competitive attitude	ordinal	The participants individual competitive attitude	
nationality	nominal	The participant's nationality	
education	nominal	The participant's education level	
children	binary	Indicating whether the participant has children or not	
friends	ratio	The participants number of friends	

Table B.3 Variables Collected from the Survey

B.2 Instructions

Information that are only presented in Info are highlighted in *italics*.

B.2.1 Managers

As a manager, you are connected to a group of three employees which consists of anonymous participants in this study. The participants are randomly selected [Company] employees without management responsibility. The combination into groups of 3 occurs randomly. Your and your group's payouts depend on your and the group members' decisions. In addition, your decisions determine the payouts of up to six additional groups.

Decision-making situation of the group members
Each member of the group must decide on the use of 10 tokens each. You and the other group members can put the 10 tokens into a private account, or you can deposit them in whole or in part into a common account. Any tokens that you do not deposit into the common account are automatically added to your private account.

Income of the group members
The total income of a group member is the sum of income from his/her private account and his/her income from the common account:

- Income from the private account: He/she earns exactly one euro for each token he/she puts in his/her private account. For example, if he/she put 4 tokens into the private account, he/she will earn exactly €4 from the private account. No one but he/she receives income from his/her private account.
- Income from the common account: For each token that is added to the common account, each group member will receive €0.5. I.e., the other two group members also each receive €0.5 for each token contribute. Conversely, the contributing group member also earns money from the contributions of the other two group members to the common account.

Your income
You as a manager will receive €15 for your participation. In addition to this €15, you also receives €0.50 for each token that your group members contribute to the shared account.

You do not earn from the deposits of your group members into the private accounts.

Your Decision

Before your group members make the contribution decisions, you decide whether or not to pay the group member with the highest contribution to the common account an additional payment of €3 to his / her private account. In the event of a tie, the €3 will be divided among all group members with the same contribution to the common account. If you opt for this additional payment scheme, this will cost you €5. If you decide against this, you will not incur any costs and no additional payments will be made to the group members.

What do the group members know about your decision?

Before making any decisions, all group members will be informed that you, the manager, decide on the additional payment of €3. Your group members also know that the additional payment is costly for you and that you earn from the deposits into the community account.

Tip for you as a manager

369 employees have already made their decision to allocate the 10 tokens between the private account and the common account. There was no additional payment for these decisions in place. On average, 2.10 Tokens were paid into the private account and 7.90 Tokens into the common account.

Summary

- All group members decide how many of the 10 tokens they deposit into their private account and how many of the 10 tokens they deposit into the common account.
- Each group member earns one euro for the tokens in the respective private account and €0.50 for each contributed token in the common account.
- You as a manager earn €0.50 for each token contributed in the common account. You cannot contribute tokens to the community account.
- The manager knows the average contribution of 369 other [Company] employees to the common account. There was no additional payment in place for these decisions.

- As a manager, you have to decide whether to pay the group member with the highest contribution to the common account an additional payment of €3 to their private. The additional payment will cost you €5.
- *In decision-making situations without additional payment, 396 [Company] employees paid an average of 2.10 tokens in the private account and 7.90 tokens in the common account.*

Comprehension Questions

Please answer the following questions to ensure that you have understood the instructions for Part I of the experiment. If you are unsure, you can return to the instructions by clicking on "Back".

Assume that none of the group members pay a contribution into the group account.

- What is the total income (private account + common account) of a group member in tokens?
- What is your income from the group's common account in euros?

Assume that all three group members each pay a contribution of 10 tokens into the group account.

- What is the total income (private account + common account) of a group member in tokens?
- What is your income from the group's common account in euros?

Assume that in a group, member A pays 0 tokens to the shared account, member B 5 tokens, and member C 10 tokens. Which member receives the additional payment of 3 tokens if the manager has selected this scheme? Member A / Member B / or Member C

Incentive Choice and Belief Elicitations

Please choose whether you want to pay the member with the highest contribution to the common account the additional payment of 3 €to his / her private account. This additional payment will cost you €5. Yes. The additional payment is used. / No. The additional payment is not used.

In addition to your earnings from your private and common account, you will receive a further payout for estimating the average contribution of the other two members of your group to your common account. Your payout will depend on how accurately you estimate the actual average contribution of your two group members. If you are exactly right, you will receive an additional €2.5 for each correct answer. If your estimate differs by 0.5 or more tokens from the actual average contribution, you will receive €0. Please enter a number from 0 to 10 (each number is allowed in steps of 0.5).

- What do you think is the average contribution of your group members' tokens to the common account with additional payment?
- What do you think is the average contribution of your group members' tokens to the common account without additional payment?
- What is the average expectation of the group members about the contribution of the other group members to the common account with additional payment?
- What is the average expectation of the group members about the contribution of the other group members to the common account with additional payment?

B.2.2 Employees

You are a member of a group of three, consisting of anonymous participants in this study. All participants are randomly selected employees of [Company]. The combination into groups of 3 occurs randomly. Your group will be connected to a manager. The manager is a randomly selected [Company] manager, i.e. a [Company] employee with management responsibility.The payouts for you, and the other group members and your manager in this section depend on your decisions, and the decisions of the other members of your group, and the manager's decision.

Decision-making situation
Each member of the group must decide on the use of 10 tokens each. You and the other group members can put the 10 tokens into a private account, or you can deposit them in whole or in part into a common account. Any tokens that you do not deposit into the common account are automatically added to your private account.

Total income
Your total income is the sum of your income from your private account and your income from the common account:

- Income from the private account: You earn exactly one euro for each token you put in your private account. For example, if you put 4 tokens into your private account, you will earn exactly €4 from your private account. No one but you receives income from your private account.
- Income from the common account: For each token that is added to the common account, you will receive €0.5. The other two group members also each receive €0.5 for each token you contribute. Conversely, you also earn money from the contributions of the other two group members to the common account. For example, if the sum of all three group members' contributions to the common account results in 30 tokens, then you and the other two group members each receive 30 x 0.5 = €15 from the common account. If the three group members pay a total of 10 tokens into the common account, you and the other two group members receive 10 x 0.5 = €5 each from the common account.

Income of you manager
Your manager will receive €15 for his / her participation. In addition to this €15, he / she also receives €0.50 for each token that you and your group members contribute to the shared account. The manager does not earn from your deposits and the deposits of your group members into the private accounts.

Decision of your manager
Before you and your group members make the contribution decisions, your manager decides whether or not to pay the group member with the highest contribution to the common account an additional payment of €3 to his / her private account. In the event of a tie, the €3 will be divided among all group members with the same contribution to the common account. If your manager decides on the additional payment, this costs the manager €5. If he / she decides against this, the manager incurs no costs and no additional payments are made to the group members.

What does the manager know when making a decision?
The manager received information about the average contribution decision of 369 other employees. These employees have already decided on the allocation of the 10 tokens between

the private account and the common account. There was no additional payment for these decisions in place. The manager *also* knows your decision-making situation. So he / she knows how much you earn, what your decision looks like and he / she also knows that you know about his / her decision. The manager doesn't know how much you and your group members are contributing when taking his/her decision on the additional payment.

Your entries

As described above, you can use 10 tokens to fund your private account and the common account. Each group member has to make two types of contribution decisions, which we will refer to below as the contribution and the contribution table. You can find a detailed description of your entries on the entry screens. When you make your decisions, you do not yet know whether the manager has selected the additional payment or not. That is why you make every decision for both scenarios - once with and once without additional payment. Since both scenarios can be relevant to your payout, you should think carefully about your decisions in both scenarios.

Summary

- All group members decide how many of the 10 tokens they deposit into your private account and how many of the 10 tokens they deposit into the common account.
- Each group member earns one euro for the tokens in the respective private account and €0.50 for each contributed token in the common account.
- The manager also earns €0.50 for each token contributed in the common account. He / she cannot contribute tokens to the community account.
- *The manager knows the average contribution of 369 other [Company] employees to the common account. There was no additional payment in place for these decisions.*
- Before you take your decisions, your manager must decide whether he / she pays the group member with the highest contribution to the common account an additional payment of €3 to the private account or whether he / she does not pay any additional payment. The additional payment costs the manager €5.
- You do not yet know how your manager decides and make your apportionment decision in the event that he / she pays the additional payment and in the event that he / she does not pay any.

Comprehension Questions

Please answer the following questions to ensure that you have understood the instructions of the experiment. If you are unsure, you can return to the instructions by clicking on "Back". When talking about your total income, please think of the sum of the income from the private account and the common account without the possible additional payment.

1. Assume that none of the group members (even you yourself) pay a contribution into the group account.
   - How high is your total income?
   - How high is the respective total income of the other two group members?
2. Assume that all three group members (also you yourself) each pay a contribution of 10 tokens into the group account.
   - How high is your total income?
   - How high is the respective total income of the other two group members?
3. Assume that you deposit 0 tokens into the common account and that the other two members of your group deposit 10 tokens each.
   - How high is your total income?
   - How high is the respective total income of the other two group members?
4. Assume that you pay 10 tokens into the common account and the other two members of your group each pay 0 tokens.
   - How high is your total income?
   - How high is the respective total income of the other two group members?

Assume that in a group, member A pays 0 tokens to the shared account, member B 5 tokens, and member C 10 tokens. Which member receives the additional payment of 3 tokens if the manager has selected this scheme? Member A / Member B / Member C

Is the additional payment scheme costly for the manager? Yes. The manager incurs costs of €5. / No. The manager incurs no costs.

Is your manager informed about other [Company] employees' contributions before making a decision on the additional payment? Yes. / No.

Contribution Decisions

When choosing the contribution to the common account, you determine how many of the 10 tokens you want to deposit into the common account. The deposit to your private account is automatically the difference between 10 tokens and your contribution to the common account.

Please enter the amount you would like to pay into the common account (any whole-number value between and including 0 and 10 is possible), if ...

- ... the manager has not selected the additional payment
- ... the manager has selected the additional payment

Now you will be asked to fill in a contribution table. In the contribution table, you should specify how many tokens you want to pay into the common account for each possible (rounded) average contribution of the other two group members to the common account. So, depending on how much the others contribute on average, you must define your own contribution decision. For each average contribution of the other two group members, please indicate the amount you would like to pay into the common account (any whole-number value between and including 0 and 10 is possible; of course, you can also enter the same amount several times):

What is your contribution to the common account if the manager has not selected the additional payment and ...

- ... the other two group members deposit an average of 0 tokens.
- ... the other two group members deposit an average of 1 tokens.
- ... the other two group members deposit an average of 2 tokens.
- ... the other two group members deposit an average of 3 tokens.
- ... the other two group members deposit an average of 4 tokens.
- ... the other two group members deposit an average of 5 tokens.

- ... the other two group members deposit an average of 6 tokens.
- ... the other two group members deposit an average of 7 tokens.
- ... the other two group members deposit an average of 8 tokens.
- ... the other two group members deposit an average of 9 tokens.
- ... the other two group members deposit an average of 10 tokens.

What is your contribution to the common account if the manager has selected the additional payment and ...

- ... the other two group members deposit an average of 0 tokens.
- ... the other two group members deposit an average of 1 tokens.
- ... the other two group members deposit an average of 2 tokens.
- ... the other two group members deposit an average of 3 tokens.
- ... the other two group members deposit an average of 4 tokens.
- ... the other two group members deposit an average of 5 tokens.
- ... the other two group members deposit an average of 6 tokens.
- ... the other two group members deposit an average of 7 tokens.
- ... the other two group members deposit an average of 8 tokens.
- ... the other two group members deposit an average of 9 tokens.
- ... the other two group members deposit an average of 10 tokens.

Help option: The numbers in the left column are the possible (rounded) average contributions of the other two group members to the common account. You now have to specify how many tokens you want to deposit into the common account for each slider, provided that the others contribute the specified amount on average. You have to make an entry in each field. For example, you are to specify how much you contribute to the common account if the other group members deposit an average of 0 tokens into the common account; how many tokens you contribute if the others contribute an average of 1 token or 2 tokens or 3 tokens, and so on. You can enter any whole-number contribution from 0 tokens to 10 tokens in each field and, of course, the same amount several times.

Incentive Compatibility

Payout relevance of your decisions
After all study participants have made their decisions, one member is randomly selected in each group of 3. For the randomly selected member, only the contribution table filled in by him/her is relevant for decision making and payout. For the other two group members who have not been selected, only the contribution is relevant for decision-making and payout. The average of the two contributions (rounded to the next whole number) then determines the relevant conditional contribution from the third member's contribution table. Of course, you do not yet know which of your contribution decisions will be randomly selected. You must therefore carefully consider both types of contribution decisions, as both can become relevant to you.

The following graphic (Figure B.1) is intended to visualize the decision-making situation. For the randomly selected person on the right, the conditional contribution from the contribution table is relevant. For the other two group members, the contribution is relevant for payout.

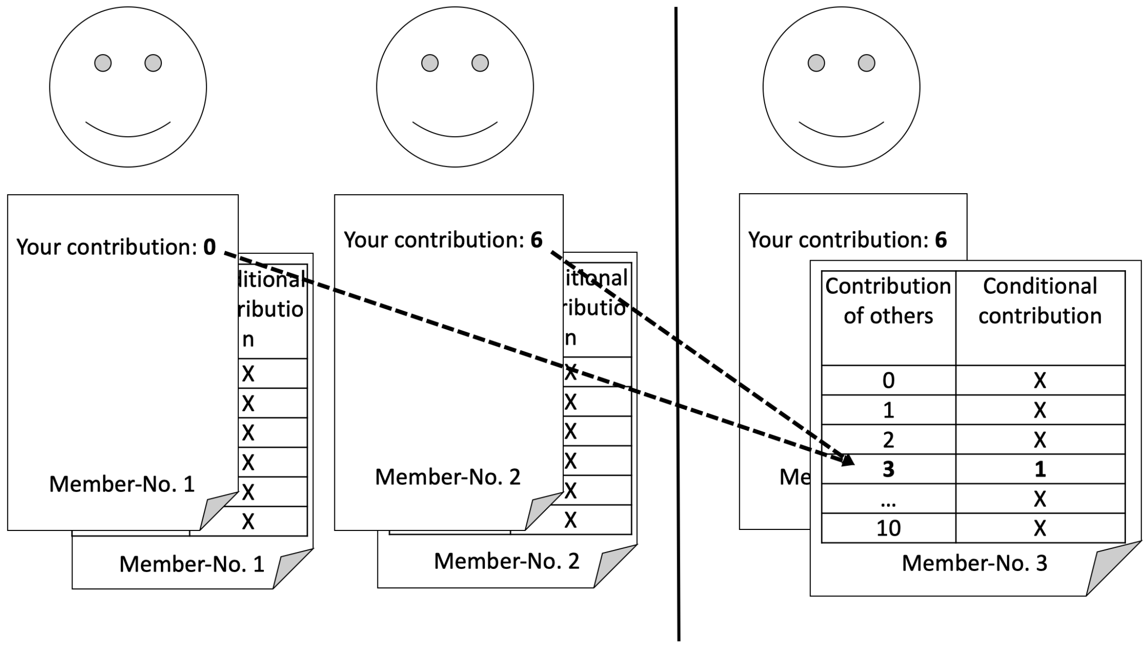

Figure B.1 Incentive Compatibility

Belief Elicitation

In addition to your earnings from your private and common account, you will receive a further payout for estimating the average contribution of the other two members of your group to your common account. Your payout will depend on how accurately you estimate the actual average contribution of your two group members. If you are exactly right, you will receive an additional €5. If your estimate differs by 0.5 or more tokens from the actual average contribution, you will receive €0. Please enter a number from 0 to 10 (each number is allowed in steps of 0.5).

What do you think is the average amount of tokens your two group members contribute to the common account?

- If the manager has selected the additional payment: ...
- If the manager has not selected the additional payment: ...

What percentage of managers chooses the additional payment scheme? Please enter a number from 0% to 100% in steps of 5% points. If you are exactly right, you will receive €1.50. If your estimate is 5 percentage points or more away from the actual average value, you will receive €0.

Please enter a number from 0 to 10 for each of the next question (any number in steps of 0.5 is allowed). If you are exactly right, you will receive €1.00 each. If your estimate is 0.5 points or more away from the actual average value, you will receive €0.

What is the average expectation of the managers about the contribution of the group members to the common account if ...

- ... the manager has not selected the additional payment
- ... the manager has selected the additional payment

B.3 Balance Tables

	Info	No Info	P-Value
Age	44.83 (10.85)	43.13 (9.39)	0.678
Female	0.30 (0.47)	0.50 (0.51)	0.143
Seniority	12.31 (7.59)	11.17 (6.42)	0.523
Career Level			
Low	0.04 (0.20)	0.04 (0.21)	1.000
Medium	0.88 (0.28)	0.87 (0.34)	1.000
High	0.08 (0.28)	0.09 (0.29)	1.000
N	23	24	

Table B.4 Balance Table Managers

Notes: P-values rely on two-sample Mann-Whitney-U tests for continuous variables or on Fisher Exact tests for categorical variables. Career levels subsume several categories in each presented category. Job functions are not shown in the table because there exists too many categories, but there are no significant differences between treatment observable. Many managers have special bonus contracts such that I do not show the variable *Individual Performance Pay* here.

	Info	No Info	P-Value
Age	36.70 (8.65)	35.57 (8.00)	0.252
Female	0.30 (0.46)	0.36 (0.48)	0.168
Seniority	4.97 (4.14)	5.19 (3.62)	0.243
Career Level			
Low	0.12 (0.33)	0.14 (0.35)	0.537
Medium	0.85 (0.36)	0.84 (0.37)	0.848
High	0.03 (0.17)	0.02 (0.12)	0.345
Indv. Perf. Pay	0.28 (0.45)	0.31 (0.47)	0.441
N	201	196	

Table B.5 Balance Table Employees

Notes: P-values rely on two-sample Mann-Whitney-U tests for continuous variables or on χ^2 -tests for categorical variables. Career levels subsume several categories in each presented category. Job functions are not shown in the table because there exists too many categories, but there are no significant differences between treatment observable.

B.4 Signaling Effects and Strategic Sophistication

	(1) Belief	(2) Belief
I(Incentive)	2.106***	2.188***
	(0.242)	(0.252)
I(Info)	-0.406	0.242
	(0.434)	(0.506)
I(Incentive)× I(Info)	0.304	0.288
	(0.394)	(0.384)
Constant	4.954***	5.474***
	(0.314)	(0.366)
Observations	406	396
R^2	0.137	0.134

Table B.6 Treatment Effects on Beliefs by Self-Evaluation of Math Skills

Notes: For each employee and dependent variable two entries are observed: one entry under the incentive and one without the incentive. Standard errors are clustered on the subject level and are shown in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix C

Spin Doctors: An Experiment on Vague Disclosure

C.1 Proofs

C.1.1 Proof of Proposition 3

Prediction 1a and 1b follow straightforwardly from proposition 1 and 2. Next, let Eg^{flex} and Eg^{prec} denote R's expected guess under flexible and precise language

$$Eg^{flex} = \chi Eg_{\chi}^{flex} + (1-\chi)Eg_{1-\chi}^{flex}$$
$$Eg^{prec} = \chi Eg_{\chi}^{prec} + (1-\chi)Eg_{1-\chi}^{prec},$$

where the subscripts χ and $1-\chi$ refer respectively to the expected guess of a naive and a sophisticated R. Also, let $\mu = 1/2$ represent the prior mean.

By the law of iterated expectations, $Eg_{1-\chi}^{flex} = Eg_{1-\chi}^{prec} = \mu$. Thus, proving prediction 2a boils down to show that $Eg_{\chi}^{flex} > Eg_{\chi}^{prec} > \mu$. Let us denote by $g_{\chi}^{flex}(\omega)$ and $g_{\chi}^{prec}(\omega)$ the equilibrium guess of a naive R when the state is ω under precise and flexible language, respectively. For any ω, $g_{\chi}^{flex}(\omega) \geq g_{\chi}^{prec}(\omega)$, with strict inequality unless $\omega = 0$ or $\omega = 1$, so that $Eg_{\chi}^{flex} > Eg_{\chi}^{prec}$. Moreover, under precise language $g_{\chi}^{flex}(\omega) \geq \omega$ with strict inequality whenever $\omega \leq \omega^*$, so that $Eg_{\chi}^{prec} > \mu$.

As for prediction 2b, Eg^{flex} is strictly increasing in χ since Eg_{χ}^{flex} and $Eg_{1-\chi}^{flex}$ are independent from χ and $Eg_{\chi}^{flex} > Eg_{1-\chi}^{flex}$. Eg^{prec} is strictly increasing in χ since $Eg_{\chi}^{prec} > Eg_{1-\chi}^{prec}$ and, moreover, while $Eg_{1-\chi}^{prec}$ is independent from χ, $Eg_{1-\chi}^{prec}$ is strictly increasing. Indeed

$$Eg_{\chi}^{prec} = \int_0^{\omega^*} \mu \, \mathrm{d}\omega + \int_{\omega^*}^1 \omega \, \mathrm{d}\omega = \mu + \int_0^{\omega^*} (\mu - \omega) \, \mathrm{d}\omega,$$

which is strictly increasing since so is ω^* and, as $\omega^* < \mu$, $\mu - \omega > 0$ in the relevant integration range.

C.1.2 Proof of Proposition 4

Prediction 3a follows directly from prediction 2a. Using the notation introduced in the previous proof (section C.1.1), prediction 3b follows from the fact that for types $\omega > \omega^*$, $g_{\chi}^{flex}(\omega) > g_{\chi}^{prec}(\omega)$ and $g_{1-\chi}^{flex}(\omega) = g_{1-\chi}^{prec}(\omega)$, while for types $\omega \leq \omega^*$ precise language is preferable if and only

$$\chi\mu + (1-\chi)\frac{\omega^*}{2} \geq \chi\frac{\omega+1}{2} + (1-\chi)\omega. \tag{C.1}$$

The inequality is violated at $\omega = \omega^*/2$, verified strictly at $\omega = 0$, and, since the left-hand side and right-hand side are respectively independent from ω and strictly decreasing, there is a unique $\hat{\omega} \in (0, \omega^*/2)$ such that the inequality holds if and only if $\omega \leq \hat{\omega}$.

Next, if we denote the expected loss of R under flexible and precise language as EL^{flex} and EL^{prec}, respectively, and we use the subscript χ and $1-\chi$ to denote the expected loss of a naive and a sophisticated R, respectively, we have

$$EL^{prec} = \chi \underbrace{\int_0^{\omega^*} (\mu - \omega)^2 \, \mathrm{d}\omega}_{EL^{prec}_{\chi}} + (1-\chi) \underbrace{\int_0^{\omega^*} \left(\frac{\omega^*}{2} - \omega\right)^2 \mathrm{d}\omega}_{EL^{prec}_{1-\chi}} = \frac{\chi^2(3+\chi)}{12(1+\chi)^3}$$

$$EL^{flex} = \chi \underbrace{\int_0^{1} \left(\frac{\omega+1}{2} - \omega\right)^2 \mathrm{d}\omega}_{EL^{flex}_{\chi}} + (1-\chi) \underbrace{0}_{EL^{flex}_{1-\chi}} = \frac{\chi}{12}.$$

Prediction 4a follows from $EL^{prec}_{1-\chi} > EL^{flex}_{1-\chi} = 0$ and

$$EL^{flex}_{\chi} > \int_0^{\omega^*} \left(\frac{\omega+1}{2} - \omega\right)^2 \mathrm{d}\omega > \int_0^{\omega^*} (\mu - \omega)^2 \, \mathrm{d}\omega = EL^{prec}_{\chi}.$$

The last inequality holds because in the relevant integration range $\omega < \mu < \frac{\omega+1}{2}$. Finally prediction 4a follows from analytical inspection, i.e., $EL^{prec} = cEL^{flex}$ with $c = \frac{\chi(3+\chi)}{(1+\chi)^3} < 1$.

C.2 Extensions

C.2.1 General Distribution of Naivete

In this section we suppose the belief of a χ-naive R upon any given message is a mixture of the posterior of a fully sophisticated receiver (with weight $1-\chi$) and a fully naive receiver (with weight χ). Besides, we assume R's type χ is drawn form a continuous distribution $h(\chi)$ with full support on $[0,1]$, mean λ and variance σ^2. The binary model we use in section 3.3.2 hence obtains as limit and special case when $h(\chi)$ puts weight only on 0 and 1. Likewise, the model in Eyster and Rabin (2005) corresponds to a degenerate $h(\chi)$ that puts all weight on a single value of χ.

One can easily verify that proposition 1 still describes S's behavior under flexible language, so that the guess of a χ-naive R upon message $[a,b]$ with $b \geq a$ is $g_\chi\big([a,b]\big) = \chi(a+b)/2 + (1-\chi)a$. As for S's behavior under precise language, it is as in proposition 2 except that S's disclosure cutoff now must solve

$$\omega^* = \int_0^1 \left(\chi\frac{1}{2} + (1-\chi)\frac{\omega^*}{2}\right) h(\chi)\mathrm{d}\chi.$$

The unique solution is $\omega^* = \frac{\lambda}{1+\lambda}$. The guess of χ-naive R upon nondisclosure is then $\chi\frac{1}{2} + (1-\chi)\frac{\omega^*}{2}$.

Thus, all results of section 3.3.2 generalize to this more flexible model. In particular, we formally establish an equivalent of predictions 4a and 4b.

Proposition 5. *For any distribution of naivete in the population*

B.a the ex-ante expected payoff of R is higher under precise language than under flexible language

B.b ex-post, the expected payoff of a χ-naive R is higher under precise language than under flexible language if and only if χ is above some cutoff $\chi^* \in (0,1)$.

Proof. The expected loss of R under flexible and precise language are now

$$EL^{flex} = \int_0^1 \int_0^1 \left(\chi \frac{\omega+1}{2} + (1-\chi)\omega - \omega\right)^2 \mathrm{d}\omega\, h(\chi)\mathrm{d}\chi = \int_0^1 \underbrace{\frac{\chi^2}{12}}_{EL_\chi^{flex}(\chi)} h(\chi)\mathrm{d}\chi$$

$$\begin{aligned} EL^{prec} &= \int_0^1 \int_0^{\omega^*} \left(\chi \frac{1}{2} + (1-\chi)\frac{\omega^*}{2} - \omega\right)^2 \mathrm{d}\omega\, h(\chi)\mathrm{d}\chi \\ &= \int_0^1 \underbrace{\frac{1}{12}\omega^* \left((\omega^*)^2 + 3(1-\omega^*)^2\chi^2\right)}_{EL_\chi^{prec}(\chi)} h(\chi)\mathrm{d}\chi, \end{aligned}$$

where EL_χ^{flex} and EL_χ^{prec} denote the expected loss of type χ. Prediction B.b follows from

$$\begin{aligned} EL_\chi^{flex}(0) &= 0 < \frac{(\omega^*)^3}{12} = EL_\chi^{prec}(0), \\ EL_\chi^{flex}(1) &= \frac{1}{12} > \frac{1}{2}\omega^*(3 - 6\omega^* + 4(\omega^*)^2) = EL_\chi^{prec}(1), \\ \frac{dEa^{flex}}{d\chi} &= \frac{\chi}{6} > \frac{1}{2}\omega^*(1-\omega^*)^2\chi = \frac{dEa^{prec}}{d\chi}. \end{aligned}$$

As for prediction B.a, we may write

$$\begin{aligned} EL^{prec} - EL^{flex} &= \frac{1}{12}\int_0^1 \left((\omega^*)^3 - (1 - 3(1-\omega^*)^2\omega^*)\chi^2\right) h(\chi)\mathrm{d}\chi \\ &\propto \frac{\lambda^3 - (1+\lambda^2(3+\lambda))\mathbb{E}\left[\chi^2\right]}{(1+\lambda)^3}. \end{aligned}$$

Thus, $EL^{prec} \geq EL^{flex}$ if and only if

$$\mathbb{E}\left[\chi^2\right] \leq \frac{\lambda^3}{1+3\lambda^2+\lambda^3}.$$

Using $\mathbb{E}\left[\chi^2\right] \equiv \lambda^2 + \sigma^2$, one sees that this is impossible as $\sigma^2 > 0$ and $\lambda^2 > \frac{\lambda^3}{1+3\lambda^2+\lambda^3}$. □

C.2.2 General Distribution of the State of Nature

The equilibrium behavior described at proposition 1 and 2 naturally generalizes to any arbitrary prior distribution $f(\omega)$ which is continuous and has full-support in the interior of $[0,1]$. Let $F(\omega)$ denote its cumulative distribution and μ the prior mean. Under flexible language, the equilibrium is identical except that the guess of a naive R upon message

$[a,b]$ with $b \geq a$ is now

$$g_\chi\big([a,b]\big) = \mathbb{E}\big[\omega \mid \omega \in [a,b]\big] = \frac{\int_a^b \omega f(\omega)\,\mathrm{d}\omega}{F(b)-F(a)}.$$

Under precise language, the equilibrium is again characterized by a disclosure cutoff $\omega^* \in (0,1)$. For a given ω^*, the guess of a rational and naive R upon nondisclosure are now respectively $g_{1-\chi}(\emptyset) = \mathbb{E}\left[\omega \mid \omega < \omega^*\right] = \frac{\int_0^{\omega^*} \omega f(\omega)\mathrm{d}\omega}{F(\omega^*)}$ and $g_\chi(\emptyset) = \mu$, so that the disclosure cutoff now solves[1]

$$\omega^* = \chi\mu + (1-\chi)\frac{\int_0^{\omega^*} \omega f(\omega)\,\mathrm{d}\omega}{F(\omega^*)}. \tag{C.2}$$

The expected loss of R under flexible and precise language are then

$$EL^{flex} = \underbrace{\chi \int_0^1 \left(\frac{\int_\omega^1 t f(t)\,\mathrm{d}t}{1-F(\omega)} - \omega\right)^2 f(\omega)\mathrm{d}\omega}_{EL_\chi^{flex}}$$

$$EL^{prec} = \underbrace{\chi \int_0^{\omega^*} \left(\int_0^1 t f(t)\,\mathrm{d}t - \omega\right)^2 f(\omega)\,\mathrm{d}\omega}_{EL_\chi^{prec}}$$

$$+ (1-\chi)\underbrace{\int_0^{\omega^*} \left(\frac{\int_0^{\omega^*} t f(t)\,\mathrm{d}t}{F(\omega^*)} - \omega\right)^2 f(\omega)\,\mathrm{d}\omega}_{EL_{1-\chi}^{prec}}.$$

All predictions of section 3.3.2 other than 4a easily extend to this setting and their proofs at section C.1 intentionally rely on general arguments.[2] In particular, prediction 4b obtains since $EL_{1-\chi}^{prec} > EL_{1-\chi}^{flex} = 0$ and

$$EL_\chi^{flex} > \int_0^{\omega^*} \left(\mathbb{E}\big[\omega \mid \omega \in [\omega,1]\big] - \omega\right)^2 f(\omega)\,\mathrm{d}\omega > \int_0^{\omega^*} (\mu-\omega)^2 f(\omega)\,\mathrm{d}\omega = EL_\chi^{prec},$$

[1] The solution is not necessarily unique. A sufficient condition for this to be the case is that $f(\omega)$ is log-concave.

[2] For the sake of precision, prediction 2b now requires equation C.2 to have a unique solution and prediction 3b requires the equivalent of inequality C.1, i.e.,

$$\chi\mu + (1-\chi)\mathbb{E}\left[\omega \mid \omega < \omega^*\right] \geq \chi\mathbb{E}\left[\omega \mid \omega \geq \omega^*\right] + (1-\chi)\omega,$$

to hold with equality for a unique ω. Log-concavity of $f(\omega)$ guarantees both.

where again the last inequality holds since in the relevant integration range $\omega < \mu < \mathbb{E}\left[\omega \mid \omega \in [\omega, 1]\right]$.

As for prediction 4a, we investigate it by simulation. In detail, we use as family of priors the beta distribution, which is defined on $[0,1]$ and can take a wide range of shapes (u-shaped, hill-shaped, increasing, decreasing) depending on its parameters $\alpha > 0$ and $\beta > 0$.[3] We numerically solve the model for different values of α and β, each ranging from $1/10$ to 10, and of χ, ranging from $1/20$ to $19/20$, and check whether $EL^{flex} - EL^{prec} > 0$.[4]

The inequality is verified for 2226 out of 2250 parameter combinations. The 24 counterexamples all obtain when $\alpha \geq 3$ and $\beta = 1/10$, i.e., when the prior mean ($\frac{\alpha}{\alpha+\beta}$) is very large and the probability mass concentrated around 1.[5] The sender's ability to make upwardly vague claims under flexible language is then somehow limited by construction, while the set of types who disclose under precise language can be very small. This explains why information transmission may eventually be higher under flexible language. This occurs for intermediate levels of naivete in the population (in all counterexamples $\chi \in [13/20, 17/20]$), so that the disclosure cutoff under precise language remains large while at the same time the welfare of sophisticates has non-negligible weight in average receiver's welfare. Notice, however, that in all counterexamples the percentage reduction in information transmission that imposing precise language entails is small, i.e., $\frac{EL^{flex} - EL^{prec}}{EL^{flex}} < -4\%$, while in "‘regular"’ instances the correspondent percentage gain is typically larger (larger than 20% in 95% of the regular instances, and as high as 99%).[6]

[3] The density of a beta distribution with shape parameters $\alpha > 0$ and $\beta > 0$ is

$$f(\omega) = \frac{\omega^{\alpha-1}(1-\omega)^{\beta-1}}{\int_0^1 t^{\alpha-1}(1-t)^{\beta-1}\mathrm{d}t}.$$

[4] When $\alpha < 1$ or $\beta < 1$, equation (C.2) can in principle have multiple solutions. Since EL^{prec} is increasing in the disclosure cutoff, we programmed both a more stringent test which uses the largest solution and a weaker test which uses the smallest one. This precaution proved unnecessary as in all instances ω^* turned out to be unique.

[5] When $\alpha > 1$ and $\beta < 1$, the density of the beta distribution is hyperbolic increasing with a vertical asymptote at 1 and, as α/β increases, the distribution gets steeper at high values of ω and flatter elsewhere.

[6] Interestingly, the highest percentage gains from imposing precise language obtain for the same distributions that generate counterexamples but for different fractions of naives, namely, for χ very small. This suggests that imposing precise language might still be on average preferable even for these prior distributions if the regulator faces some uncertainty about the level of sophistication in the population.

C.3 Evolution of Play

After each round, each subject receives feedback about the sender's type and the receiver's guess in her pair. Tables C.1 and C.2 summarize how sender behavior evolves over rounds in Flexible and Precise, respectively. We split the total number of rounds in three phases of five rounds each, i.e phase 1 (rounds 1 to 5), phase 2 (rounds 6 to 10) and phase 3 (rounds 11 to 15). In Flexible, the most frequent messages of types 3, 4, and 5 coincide with the theoretical predictions in all phases. Types 1 and 2 in rounds 1-5 use most frequently nondisclosure rather than their theoretically predicted message, i.e., respectively $\{1,2,3,4,5\}$ and $\{2,3,4,5\}$. However, over time their behavior get closer to the theoretical predictions. In particular, in rounds 6-10 and 11-15 types 2 most frequently send the predicted message. Likewise, in rounds 11-15, the predicted message of type 1 is almost as frequent as nondisclosure (14 subjects of type 1 do not disclose, 12 subjects send $\{1,2,3,4,5\}$ and 2 subjects send $\{1,2,3,4\}$).

In Precise, we observe that over time disclosure rates generally increase for high types and decrease for low types. The sharp increase in the disclosure rate of types 2 is likely to reflect a shift in the disclosure threshold, i.e., a strategic response to the increase in receiver skepticism documented below. Instead, the increase in the disclosure rate of types 5 is likely to be the result of learning, i.e., a reduction in noisy behavior.

	type=0	type=1	type=2	type=3	type=4	type=5
rounds 1-5	nondiscl.	nondiscl.	nondiscl.	$\{3,4,5\}$	$\{4,5\}$	$\{5\}$
rounds 6-10	nondiscl.	nondiscl.	$\{2,3,4,5\}$	$\{3,4,5\}$	$\{4,5\}$	$\{5\}$
rounds 11-15	nondiscl.	nondiscl.	$\{2,3,4,5\}$	$\{3,4,5\}$	$\{4,5\}$	$\{5\}$

Table C.1 Modal Sender Messages Over Time in the Flexible Treatment

	type=0	type=1	type=2	type=3	type=4	type=5
rounds 1-5	6.3%	10%	20.6%	71.1%	83.8%	70.8%
rounds 6-10	3.5%	3%	38.9%	92.1%	93.1%	93.3%
rounds 11-15	5%	3%	66.7%	85.2%	93%	92.9%

Table C.2 Disclosure Rates Over Time in the Precise Treatment

On the receiver side, the average normalized guess decreases over time, suggesting that receivers become more skeptical (see also prediction 2b). In Flexible, the average normalized guess is 32% in rounds 1-5, 23.1% in rounds 6-10, and 19% in rounds 11-15. In Precise, it decreases from 38.2%, to 33.7% in rounds 6-10, and to 27% in rounds 11-15. The regressions in Table 3.1 show that there is also a negative time trend in receiver mistakes. However, receivers keep significantly overestimating sender types in all phases (two-tailed t-tests with clustering on subject and on pair level, for all phases p-value< 0.001).

Columns (1) to (4) in Table C.3 document no differential treatment effect on information transmission across phases. In column (5), we observe that the treatment effect for sophisticated receivers significantly increases in the direction predicted by theory in rounds 11-15. Column (6) confirms the estimation from Table 3.1 and shows no significant variation over time. However, we acknowledge that the analysis of interactions between time effects and treatment may suffer from a lack of power.

	(1) Mistake	(2) Mistake	(3) Mistake	(4) Mistake	(5) Mistake	(6) Mistake
Precise (d)	-0.115	-0.196*	-0.227*	-0.316***	-0.0267	-0.0404
	(0.0911)	(0.101)	(0.126)	(0.0993)	(0.110)	(0.108)
Rounds 1-5 (d)	0	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)	(.)
Rounds 6-10 (d)	-0.257***	-0.188**	-0.302***	-0.314***	-0.284***	-0.148
	(0.0692)	(0.0757)	(0.107)	(0.0990)	(0.0768)	(0.0919)
Rounds 11-15 (d)	-0.349***	-0.276***	-0.328**	-0.434***	-0.380***	-0.128
	(0.0785)	(0.0922)	(0.130)	(0.142)	(0.0841)	(0.109)
Precise × Rounds 1-5 (d)	0	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)	(.)
Precise × Rounds 6-10 (d)	0.0694	0.0959	0.0247	0.113	0.209	0.177
	(0.107)	(0.0989)	(0.160)	(0.134)	(0.132)	(0.113)
Precise × Rounds 11-15 (d)	0.0911	0.0641	-0.103	0.148	0.292**	0.0423
	(0.124)	(0.111)	(0.186)	(0.173)	(0.141)	(0.118)
Constant	1.624***	1.757***	2.120***	2.319***	1.270***	1.299***
	(0.102)	(0.119)	(0.165)	(0.156)	(0.113)	(0.133)
Type dummies	Yes	Yes	Yes	Yes	Yes	Yes
Incl. sender choices	All	Theory-conforming	All	Theory-conforming	All	Theory-conforming
Incl. receivers	All	All	Naives	Naives	Soph.	Soph.
R^2	0.173	0.457	0.338	0.608	0.136	0.454
Observations	1185	828	510	360	675	468

Table C.3 Regressions on Receivers' Absolute Mistakes Over Time

Notes: Robust standard errors clustered at the subject level in parentheses; $^{*}p < 0.10,^{**}p < 0.05,^{***}p < 0.01$.

C.4 Alternative Classification of Naives and Sophisticates

This appendix demonstrates the robustness of results in columns (3) through (6) of Table 3.1 by repeating the analysis using different classifications of sophistication and naivete. In columns (1) and (2) of Table C.4, receivers are classified as sophisticated if they are "rarely naive", i.e., if they make the fully naive choice in less than 30 percent of the rounds in which they face either vague disclosure or nondisclosure. In columns (3) and (4) of Table C.4, receivers are classified as sophisticated if they are "never naive", i.e., if they never make a fully naive choice in the rounds in which they face either vague disclosure or nondisclosure. In columns (5) and (6) of Table C.4, we use a measure that is exogenous to receiver's choices in the experiment for the classification: we classify receivers with a high school math grade (Abitur) of 1 or 2 as sophisticated and receivers with a math grade of 3, 4, 5 or 6 as naives. This classification is equivalent to a median split.

When we use the "rarely naive" criterion, we find that naives make significantly smaller mistakes under precise language and that sophisticates make insignificantly larger mistakes under precise language. This mimics our findings when we use the "hardly ever naive" criterion in the main text. When we use the "never naive" criterion (columns (3) and (4)), naives make insignificantly smaller mistakes under precise language, while sophisticates make significantly larger mistakes. The "never naive" criterion results in a more selective pool of sophisticates who are hurt by moving from flexible to precise language. When we classify receivers based on their high school math grade (columns (5) and (6)), we find that naives do significantly worse and that sophisticates do insignificantly better under precise language.

The above criteria, except for the math grade, are based on the incidence of naive choices and therefore pool all other choices under the label of sophisticated behavior. Alternatively, we may call a receiver sophisticated if her choices line up well with empirical best response behavior. Table C.4 uses three notions of empirical best response behavior to classify receivers. Consider the criterion "best response 1". As in section C.3, we divide our experiment into phase 1 (rounds 1 to 5), phase 2 (rounds 6 to 10) and phase 3 (rounds 11 to 15). For each phase and each possible message, including nondisclosure, we calculate the average type that actually sent this message. The use of phases allows us to arrive at a more precise measure of average behavior. We call a receiver's guess a noisy empirical best response if it lies less than 0.5 above and less than 0.5 below the average sender type conditional on a given message. The criterion "best response 1" then classifies a receiver as sophisticated if her guess is a noisy best response in more than 75 percent of rounds

	(1) Mistake	(2) Mistake	(3) Mistake	(4) Mistake	(5) Mistake	(6) Mistake
Precise (d)	-0.186*	0.0294	-0.0856	0.188***	-0.343***	0.0265
	(0.0971)	(0.0762)	(0.0853)	(0.0633)	(0.0928)	(0.0822)
Round	-0.0202	-0.0319***	-0.0333***	-0.0207**	-0.0491***	-0.0250***
	(0.0138)	(0.00607)	(0.00755)	(0.00852)	(0.0106)	(0.00653)
Constant	2.189***	1.445***	1.977***	1.008***	2.411***	1.432***
	(0.250)	(0.112)	(0.136)	(0.0961)	(0.142)	(0.112)
Type dummies	Yes	Yes	Yes	Yes	Yes	Yes
Incl. receivers	Naives	Soph.	Naives	Soph.	Naives	Soph.
Criterion	**Rarely naive**		**Never naive**		**Math grade**	
R^2	0.361	0.138	0.265	0.116	0.375	0.130
Observations	300	885	750	435	330	855

Table C.4 Regressions on Receivers' Absolute Mistakes by Naivite Classification

Notes: Robust standard errors clustered at the subject level in parentheses; $^{*}p < 0.10,^{**} p < 0.05,^{***} p < 0.01$.

that featured either vague disclosure or nondisclosure. The criterion "best response 2" is laxer and classifies an individual as sophisticated if her guess is a noisy best response in more than 50 percent of rounds that featured either vague disclosure or nondisclosure. The criterion "best response 3" is defined like "best response 1" except that is allows for a 1-unit deviation from the true average type in defining the empirical best response.

In Table C.4, columns (1) and (2) feature the criterion "best response 1" and columns (3) and (4)the same criterion, but only theory-conforming sender behavior. Columns (5) and (6) feature best response 2, whereas columns (7) and (8) feature best response 3. In all cases, naives are found to make significantly smaller mistakes under precise language, while there is no treatment effect on sophisticates. A direct implication of our model's prediction that naives are better off and sophisticates are worse off under precise language is that the negative treatment effect of imposing precise language on receiver mistakes is larger for naives, i.e., that there is a significant interaction effect between treatment and sophistication. Table C.5 confirms that the data bears out this prediction for all criteria we have introduced above.

	(1) Mistake	(2) Mistake	(3) Mistake	(4) Mistake	(5) Mistake	(6) Mistake	(7) Mistake	(8) Mistake
Precise (d)	-0.257***	0.125	-0.359***	0.0297	-0.311***	-0.0775	-0.257***	0.00789
	(0.0810)	(0.0814)	(0.0873)	(0.0482)	(0.0811)	(0.0746)	(0.0932)	(0.0439)
Round	-0.0280***	-0.0339***	-0.0198***	-0.0206***	-0.0178	-0.0361***	-0.0219***	-0.0185***
	(0.00774)	(0.00852)	(0.00670)	(0.00572)	(0.0105)	(0.00634)	(0.00762)	(0.00439)
Constant	2.050***	1.220***	2.179***	1.219***	2.240***	1.458***	2.286***	1.238***
	(0.132)	(0.115)	(0.144)	(0.111)	(0.174)	(0.110)	(0.164)	(0.0975)
Type dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Incl. receivers	Naives	Soph.	Naives	Soph.	Naives	Soph.	Naives	Soph.
Criterion	**Best response 1**		**Best response 1**		**Best response 2**		**Best response 3**	
Sender choices	All		Theory-conforming		All		All	
R^2	0.264	0.134	0.532	0.569	0.359	0.124	0.550	0.592
Observations	720	465	519	309	420	765	410	418

Table C.5 Regressions on Receivers' Absolute Mistakes by Best Response

Notes: Robust standard errors clustered at the subject level in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1) Mistake	(2) Mistake	(3) Mistake	(4) Mistake	(5) Mistake	(6) Mistake	(7) Mistake
Precise (d)	-0.255***	0.213	-0.209**	-0.0824	-0.256***	-0.326***	-0.161*
	(0.0785)	(0.157)	(0.102)	(0.0852)	(0.0814)	(0.0915)	(0.0891)
Round	-0.0299***	-0.0299***	-0.0299***	-0.0299***	-0.0299***	-0.0299***	-0.0299***
	(0.00583)	(0.00582)	(0.00582)	(0.00582)	(0.00583)	(0.00583)	(0.00583)
Hardly ever naive (d)	-0.471***						
	(0.0684)						
Precise x Hardly naive (d)	0.384***						
	(0.111)						
Math grade		0.100*					
		(0.0515)					
Precise x Math grade (d)		-0.132**					
		(0.0615)					
Rarely naive (d)			-0.366***				
			(0.0774)				
Precise x Rarely naive (d)			0.229*				
			(0.127)				
Never naive (d)				-0.419***			
				(0.0666)			
Precise x Never naive (d)				0.247**			
				(0.107)			
Best response 1 (d)					-0.448***		
					(0.0732)		
Precise x Best resp. 1 (d)					0.345***		
					(0.111)		
Best response 2 (d)						-0.398***	
						(0.0830)	
Precise x Best resp. 2 (d)						0.229*	
						(0.118)	
Best response 3 (d)							-0.448***
							(0.0733)
Precise x Best resp. 3 (d)							0.205*
							(0.108)
Constant	1.907***	1.450***	1.908***	1.765***	1.888***	1.995***	1.892***
	(0.0978)	(0.148)	(0.100)	(0.105)	(0.105)	(0.102)	(0.106)
R^2	0.208	0.177	0.192	0.197	0.205	0.189	0.213
Observations	1185	1185	1185	1185	1185	1185	1185

Table C.6 Robustness Check on Receivers' Absolute Mistakes

Notes: Robust standard errors clustered at the subject level in parentheses; $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$.

C.5 Senders' Out-Of-Sample Beliefs

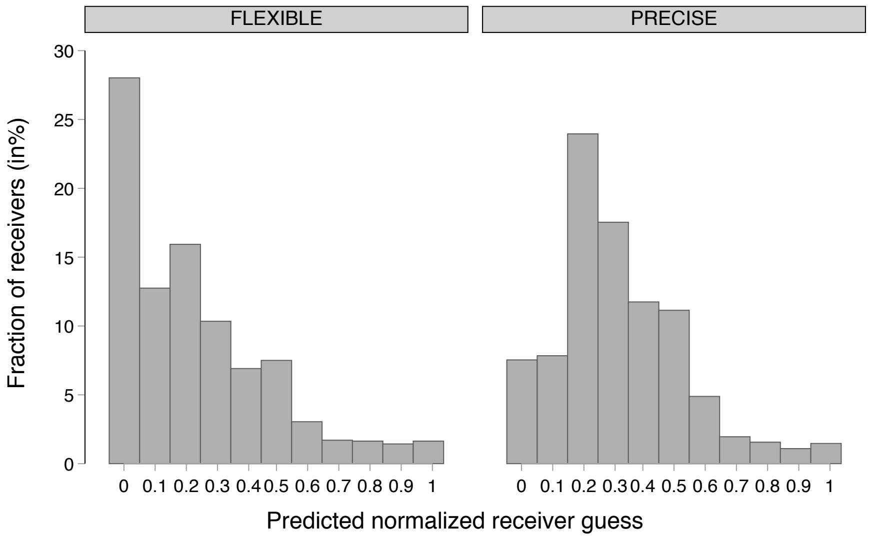

Figure C.1 Senders' Out-Of-Sample Predictions

Notes: Senders' out-of-sample predictions of receivers' distribution of guesses conditional on observing non-disclosure (by treatment). Senders made predictions about receiver behavior in a pilot experiment.

C.6 Decision Screens

Figure C.2 shows the decision screen in Flexible. The sender could freely specify the interval to send by clicking on and herewith selecting the respective types to be included. A preview window showed how the message would appear on the receiver's screen.

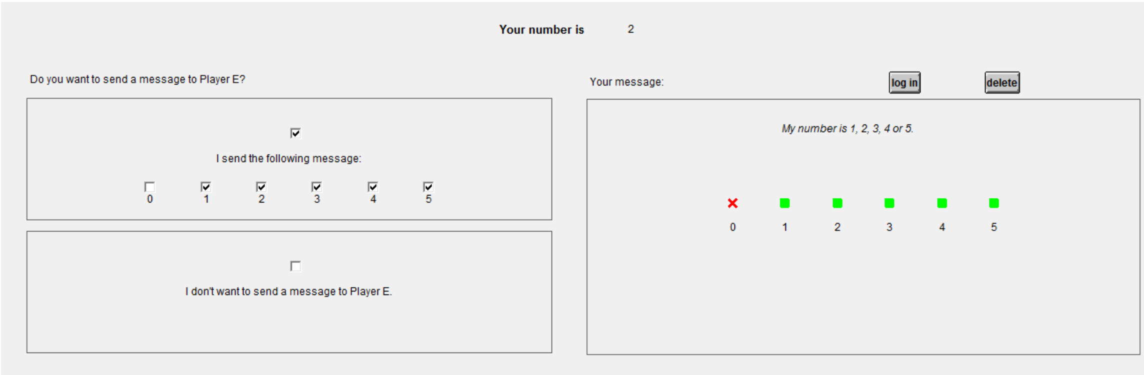

Figure C.2 Senders' Decision Screen in the Flexible Treatment

Figure C.3 shows an example if a sender decision screen in Precise. Here, the senders were provided with the two options in random order.

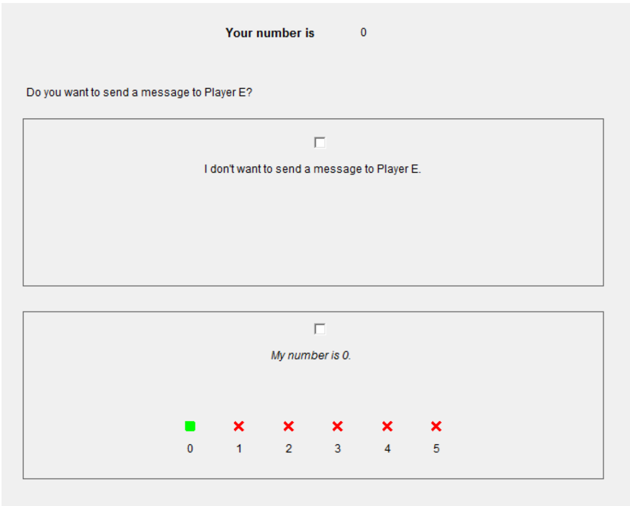

Figure C.3 Senders' Decision Screen in the Precise Treatment

C.7 Instructions

C.7.1 Flexible Treatment

This experiment is composed of 15 rounds. At the beginning of this experiment, it will be determined randomly whether you are **player S** or **player E**. You will keep this role in all 15 rounds. In each round you play a game with a randomly chosen participant in the opposite role. It is very unlikely, that you are paired up with the same participant in two consecutive rounds.

The Game
In each round, player S receives a number on the **range 0, 1, 2, 3, 4, 5** via the computer. All the numbers are equally likely. Player E does not see which number player S receives. However, player S can send a message regarding his or her number to player E. Player E must guess the number of player S. At the end of each round both players are informed about the number of player S and the guess of player E.

Decision of player S
After receiving the number, player S can decide about whether or not he or she would like to send a message to the recipient. Player S can decide which message he or she would like to send. In doing so, three rules must be complied with:

1. **The sent message must contain the true number of the sender**

   Example: If the sender receives number 3, he can only send messages that contain the number 3.

2. **The sent message must not contain gaps.**

   Example: The sender with number 3 must not send the numbers 2, 3, 5 as possible numbers because the 4 is missing in this row.

3. **The send message may contain maximum five numbers.**

   Example: The sender with the number 3 may only send 5 of is possible numbers in total. The sender may not send all six numbers (i.e., 0, 1, 2, 3, 4, 5).

When player S has received, for example, the number 3, he or she can send a message that contains the true number and no gaps or send no message at all. This, for example, applies to the message “My number is 3.”. Graphically, the message “My number is 3” will be depicted by a green box above number 3 and red crosses above 0, 1, 2, 4 and 5:

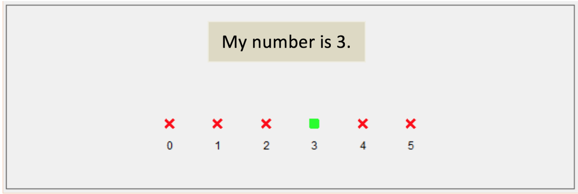

Probability	Payoff
PP%	8 Euro
(100-PP)%	1 Euro

Decision of player E

Player E either sees the message sent by player S or he or she will see the note "Player S has not sent you a message." if player S has decided not to send a message. Then, player E must enter his or her guess about the actual number of player S. Here, every number can be entered in 0.5-intervals (0, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, 5).

Payoff

The payment is determined by the following rules: The higher the guess of player E, the higher the payment of player S. And the closer the guess of player E is to the true number of player S, the higher the payment of player E.

Hereafter, the mechanism which determines the payment is explained in detail.

In each round you can earn between 0 and 100 **probability points (PP)**. The more probability points you earn, the higher the probability that you win the subsequent lottery:

If you gain 0 probability points you receive with certainty (with 100%) 1 Euro. If you gain 100 probability points you receive with certainty (with 100%) 8 Euro. If you gain e.g. 70 probability points, you receive, with the probability of 70%, 8 Euro and, with the probability of 30% 1 Euro. The more probability points you gain, the more probable it is that you receive 8 instead of 1.

Thus, you should try to gain as many probability points as possible.

The amount of your probability points in one round depends on both the number of player S and the **guess of player E**. The **payoff table**, which you can find at your spot, makes this clear. If player S e.g. receives the number 3 and player E guesses number 4.5, player E gains 79 probability points and player S 96 probability points. But, if player E guesses that the number of Player S is 1, player E gains 69 probability points and player S only 19 probability points.

Only one of the 15 rounds is chosen randomly and then is actually relevant to the payoff. Your probability points in this round determine the lottery that is played by the computer at the end of the experiment. Since you do not know, which of the 15 rounds is relevant to the payoff you should think carefully about your decisions in each round.

Summary

- Player S receives a random number that is unknown to player E.
- Player S decides whether or not to send a message to player E regarding the number. The message must contain the number of player S.
- What the message contains is determined by player S.
- Player E must guess the number of player S.
- The higher player E guesses the number of player S, the higher the chances of achieving a higher profit for player S.
- The more accurate the guess of player E for the number is, the higher the chances of profits for player E.

C.7.2 Precise Treatment

This experiment is composed of 15 rounds. At the beginning of this experiment, it will be determined randomly whether you are **player S** or **player E**. You will keep this role in all 15 rounds. In each round you play a game with a randomly chosen participant in the opposite role. It is very unlikely, that you are paired up with the same participant in two consecutive rounds.

The Game
In each round, player S receives a number on the **range 0, 1, 2, 3, 4, 5** via the computer. All the numbers are equally probable. Player E does not see which number player S receives. However, player S can send a message regarding his or her number to player E. Player E must guess the number of player S.

At the end of each round both players are informed about the number of player S and the guess of player E.

Decision of player S
After receiving the number, player S can decide about whether or not he or she would like to send a message to the recipient. If player S does send a message, player E will be informed about the number. If player S does not send a message, player E will not be informed about the number.

When player S has received e.g. the number 3, he or she can send a message that contains the true number or send no message at all. This, for example, applies to the message "My number is 3". Graphically, the message "My number is 3" will be depicted by a green box above number 3 and red crosses above 0, 1, 2, 4 and 5:

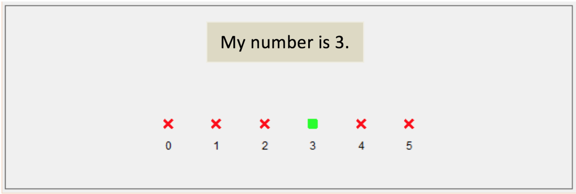

Decision of player E

Player E either sees the message sent by player S or sees the note "Player S has not sent you a message." if player S has decided not to send a message.

Then, player E must enter his or her guess about the actual number of player S. Here, every number can be entered in 0.5-intervals (0, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, 5).

Payoff

The payment is determined by the following rules: The higher the guess of player E, the higher the payment of player S. And the closer the guess of player E is to the true number of player S, the higher the payment of player E.

Hereafter, the mechanism which determines the payment is explained in detail.

In each round you can earn between 0 and 100 **probability points (PP)**. The more probability points you earn, the higher the probability that you win the subsequent lottery:

Probability	**Payoff**
PP%	8 Euro
(100-PP)%	1 Euro

If you gain 0 probability points you receive with certainty (with 100%) 1 Euro. If you gain 100 probability points you receive with certainty (with 100%) 8 Euro. If you gain e.g. 70 probability points, you receive, with the probability of 70%, 8 Euro and, with the probability of 30% 1 Euro. The more probability points you gain, the more probable it is that you receive 8 Euro instead of 1 Euro.

Thus, you should try to gain as many probability points as possible.

The amount of your probability points in one round depends on both the number of player S and the **guess of player E**. The **payoff table**, which you can find at your spot, makes this clear. If player S e.g. receives the number 3 and player E guesses number 4.5, player E gains 79 probability points and player S 96 probability points. But, if player E guesses that the number of player S is 1, then player E gains 69 probability points and player S only 19 probability points.

Only one of the 15 rounds is chosen randomly and then is actually relevant to the payoff. Your probability points in this round determine the lottery which is played by the computer at the end of the experiment. Since you do not know, which of the 15 rounds is relevant to the payoff you should think about your decisions in each round.

Summary

- Player S receives a random number that is unknown to player E.
- Player S decides whether or not to send a message to player E regarding the number. The message must contain the number of player S.
- What the message contains is determined by player S.
- Player E must guess the number of player S.
- The higher player E guesses the number of player S, the higher the chances of profits for player S.
- The more accurate the guess of player E for the number is, the higher the chances of profits for player E.

C.7.3 Payoff Tables

		Guess of Player E										
		0	**0.5**	**1**	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**
Number of Player S	**0**	100	96	88	79	69	58	46	32	19	4	0
	1	88	96	100	96	88	79	69	58	46	32	19
	2	69	79	88	96	100	96	88	79	69	58	46
	3	46	58	69	79	88	96	100	96	88	79	69
	4	19	32	46	58	69	79	88	96	100	96	88
	5	0	4	19	32	46	58	69	79	88	96	100

Table C.7 Payoffs of Player E

C.7.4 Out-Of-Sample Belief Elicitations

At the end of the experiment, subjects received the following questions based on their role and treatment condition. Some questions, as indicated below, were incentivized using the average behavior of subjects that participated in the pilot session as a benchmark.

		Guess of Player E										
		0	**0.5**	**1**	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**
Number of Player S	**0**	0	4	19	32	46	58	69	79	88	96	100
	1	0	4	19	32	46	58	69	79	88	96	100
	2	0	4	19	32	46	58	69	79	88	96	100
	3	0	4	19	32	46	58	69	79	88	96	100
	4	0	4	19	32	46	58	69	79	88	96	100
	5	0	4	19	32	46	58	69	79	88	96	100

Table C.8 Payoffs of Player S

Receiver beliefs in Precise

Please answer the following questions. You can earn additional money with your answers. In you answers, refer to the first round of an experiment that is very similar to today's experiment, but that took place with other participants, at MELESSA, several weeks ago. Your answers will be compared to the data of the previous experiment. You will receive 100 probability points (PP) (equal to a 100% chance of winning) for a lottery that gives you either €2 or €0. Then, 14 PP of the 100 PP are deducted for each incorrect answer. An input is considered incorrect, if it differs by more than 5%-points from the true value. Your input can be made without the %-sign. The sum of your inputs must be 100.

What percentage of players S who did not send a message to player E had the following number?

0:[ ]; 1:[ ]; 2:[ ]; 3:[ ]; 4:[ ]; 5:[ ]

Sender beliefs in Precise

Please answer the following questions. You can earn additional money with your answers. In you answers, refer to the first round of an experiment that is very similar to today's experiment, but that took place with other participants, at MELESSA, several weeks ago. Your answers will be compared to the data of the previous experiment. You will receive 100 probability points (PP) (equal to a 100% chance of winning) for a lottery that gives you either €2 or €0. Then, 8 PP of the 100 PP are deducted for each incorrect answer. An input is considered incorrect, if it differs by more than 5%-points from the true value. Your input can be made without the %-sign. The sum of your inputs must be 100.

What percentage of players E guessed the following number when they did not receive a message from player S?

0:[ ]; 0.5:[ ]; 1:[ ]; 1.5:[ ]; 2:[ ]; 2.5:[ ]; 3:[ ]; 3.5:[ ]; 4:[ ]; 4.5:[ ]; 5:[ ]

Receiver beliefs in Flexible

Please answer the following questions. You can earn additional money with your answers. In you answers, refer to the first round of an experiment that is very similar to today's experiment, but that took place with other participants, at MELESSA, several weeks ago. Your answers will be compared to the data of the previous experiment. You will receive 100 probability points (PP) (equal to a 100% chance of winning) for a lottery that gives you either €2 or €0. Then, 14 PP of the 100 PP are deducted for each incorrect answer. An input is considered incorrect, if it differs by more than 5%-points from the true value. Your input can be made without the %-sign. The sum of your inputs must be 100.

What percentage of players S had the following number when they did not send a message to player E?
0:[ ]; 1:[ ]; 2:[ ]; 3:[ ]; 4:[ ]; 5:[ ]

[On new screen:] Additionally, please answer the following questions. Refer again to the first round of the experiment that has already taken place at MELESSA.

What was the most common message sent to player E when player S had the following numbers?
(Please always state the upper and the lower number of a message. Example: For the message "My number is 3, 4, or 5", "3" is the lower number and "5" is the upper number. You should enter "3" in the left box and "5" in the right box. If a message only contains one number, then this number should be entered as the lower as well as the upper number. If no message is sent, leave both boxes blank.)
0:[ ] to [ ]; 1:[ ] to [ ]; 2:[ ] to [ ]; 3:[ ] to [ ]; 4:[ ] to [ ]; 5:[ ] to [ ]

Sender beliefs in Flexible

Please answer the following questions. You can earn additional money with your answers. In you answers, refer to the first round of an experiment that is very similar to today's experiment, but that took place with other participants, at MELESSA, several weeks ago. Your answers will be compared to the data of the previous experiment. You will receive 100 probability points (PP) (equal to a 100% chance of winning) for a lottery that gives you either €2 or €0. Then, 8 PP of the 100 PP are deducted for each incorrect answer. An input is considered incorrect, if it differs by more than 5%-points from the true value.

Your input can be made without the %-sign. The sum of your inputs must be 100.

What percentage of players E have guessed the following if they had not received a message from player S?
0:[ ]; 0.5:[ ]; 1:[ ]; 1.5:[ ]; 2:[ ]; 2.5:[ ]; 3:[ ]; 3.5:[ ]; 4:[ ]; 4.5:[ ]; 5:[ ]

[On new screen:] Additionally, please answer the following questions. Refer again to the first round of the experiment that has already taken place at MELESSA.

What was the average guess of player E when player S sent the following message?
"My number is 1, 2, 3, 4, or 5.":[ ] "My number is 2, 3, 4, or 5.":[ ] "My number is 3, 4, or 5.":[ ] "My number is 4 or 5.":[ ]

Appendix D

Complexity and Appropriation Interact in Affecting Compliance Behavior

D.1 Instructions

Introduction

Welcome to an experiment on decision-making behavior!
Thank you for your participation!

During the experiment, you and all other participants are asked to make decisions. Your payout will be determined according to the rules explained below.

Please do not speak with other participants of the experiment from now on. If you have any questions after the instructions or during the experiment, please press the red button on the keyboard in front of you. One of the experimenters will then come to you and answer your questions privately.

The experiment lasts a maximum of 60 minutes. All your decisions and answers remain anonymous. Neither the experimenters nor the other participants will know which decisions you have made and which participant earns how much.

All payouts from the experiment will be handed over to you privately and in cash.

Slider Task

Your task
In each of the two consecutive rounds (round A and round B), you will see 48 sliders on your screen. Your task is to bring as many of these sliders as possible to position 50:

Use both the computer mouse and the keyboard for positioning. You have 120 seconds time for one round, so to position 48 sliders.

Your payment
Each correctly positioned slider yields €0.40. The sum of all correctly positioned sliders from both rounds gives you the total amount of money that is generated in this task. However, this amount does not make your payout of this experiment. Only a share of the total amount of money is your payout. You will learn how high your share of the money

is after we have finished the task. Note, however: the more sliders you position correctly, the higher your payout of the experiment will be.

Practice round
Before you start your task, there will be two practice rounds. You will not receive any money from the two practice rounds, but you can get to know your task.

[After the slider tasks have been processed, the experimenter pays out the money earned in the slider tasks in random order.]

Compliance Decision

Your share
The share of the money you have earned from your task will be determined by the form(s) *[depending on the treatment]* you received along with these instructions. Put your share of the money in the envelope labeled “Your Share”.

The remaining share
The remaining share of the money you have earned from your task can also be determined using the form(s) *[depending on the treatment]*. Put the remaining share in the envelope labeled “Remaining Share”.

Moral Treatment

Usage of the remaining share money
The remaining share money will be used by the lab researchers. The researchers will donate the money to the German Bone Marrow Donation Registry (DKMS). The main activity of the DKMS is to improve the healing potential of leukemia and other life-threatening diseases of the blood-forming system by supporting bone marrow donations. One major part of DKMS is the DKMS umbilical cord blood bank, which collects, processes, stores and mediates umbilical cord blood stem cell donations for newborns. (Information from www.dkms.de)

Waste Treatment

Usage of the remaining share of the money
The remaining share money will be used by the lab researchers. The researchers will donate the money to the Bavarian Yacht Club (BYC). The main activity of the BYC is to

professionally promote sailing with all its modern features and high standards. In addition, the social life outside the gates of Munich is cultivated. The BYC also has an exquisite restaurant in its Clubcasino at the Starnberger Lake. (Information from www.byc.de)

The remaining amount of money will be donated by Marvin Deversi, a member of the Chair of Behavioral Economics and Experimental Economic Research of LMU Munich, on behalf of Prof. Dr. Florian Englmaier, head of the Chair of Organizational Economics at LMU Munich, to the BYC/the DKMS. The verifying documents for the total amount of donations, including the time of today's experiment, will be posted on the White Board in front of the MELESSA laboratory in the week of April 17, 2017 to April 21, 2017. We will not post personal data.

End

Seal both envelopes using the sticky tape and then click "Quit Part I" to finish this part of the experiment. You must not open any of the two envelopes during this experiment. All your decisions remain anonymous. After the experiment is over, you take the envelope labeled "Your Share" home and leave the envelope labeled "Remaining Share" on your table. The experimenter will collect the remaining envelopes only after every participant has left the room.

As soon as you click on "Next", your earnings summary from your task will be displayed again on the screen and you will be able to start calculating your and the remaining share.

D.2 Example Screens

Figure D.1 First Round of Slider Tasks

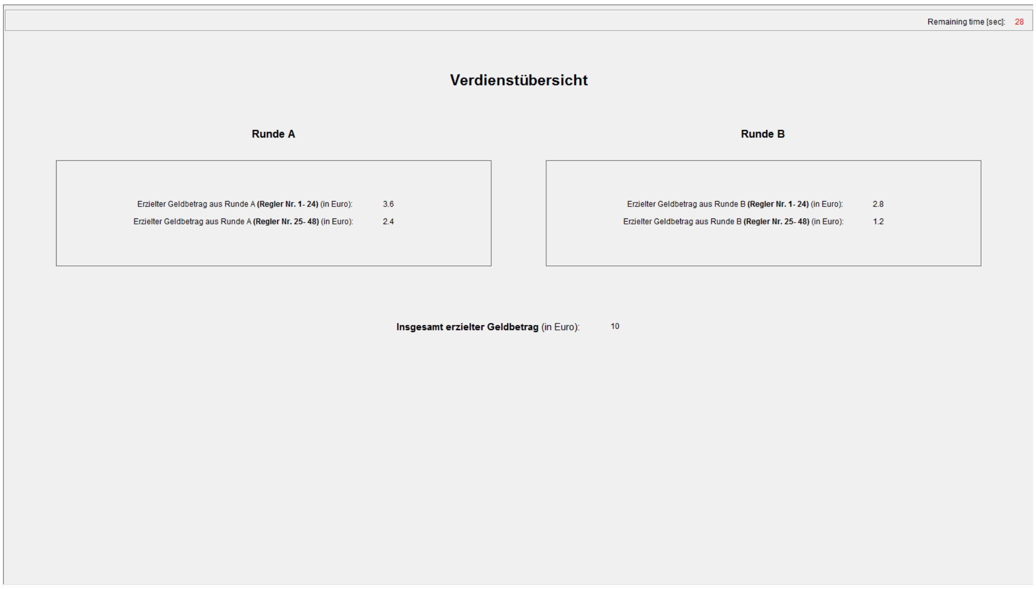

Figure D.2 Overview of Correctly Positioned Sliders

D.3 Forms

Taxable income (line 299 of your return) | 1 |

If your taxable income on line 1 **above**
- is $41,935 or less, enter it on line 2 of column **A**;
- is more than $41,935 but not more than $83,865, enter it on line 2 of column **B**;
- is more than $83,865 but not more than $102,040, enter it on line 2 of column **C**;
- is more than $102,040, enter it on line 2 of column **D**.

			A	B	C	D
Taxable income (see the instructions above)		2				
	−	3	00,000 00	41,935 00	83,865 00	102,040 00
Subtract line 3 from line 2.	=	4				
	x	5	16%	20%	24%	25.75%
Multiply line 4 by line 5.	=	6				
	+	7	00,000 00	6,709 60	15,095 60	19,457 60
Add lines 6 and 7. Carry the result to line 401 of your return. **Income tax on taxable income**	=	8				

Figure D.3 Clipping of a Tax Form in Quebéc

Notes: Example for a complex tax form in Quebéc showing that tax payers need to conduct complicated multiplications, consider if-conditions, and carry numbers across different forms.

EXAMPLE for SIMPLE

«Sitz»

Form

REMAINING SHARE

Money from slider tasks		1	22
	x		50%
Multiply row 1 with 50%. **The remaining share of the money is**	=	2	11

YOUR SHARE

Transfer amount from row 2. **Your share of the money is**	3	11

EXAMPLE for COMPLEX

Seat

Form H

REMAINING SHARE

Transfer amount from row 3 of form A.		1	6
Transfer amount from row 3 of form B.	+	2	5
Add amounts from rows 1 and 2. **Remaining share of the money is**	=	3	11

YOUR SHARE

Transfer amount from row 6 of form A.		4	6
Transfer amount from row 6 of form B.	+	5	5
Add amounts from rows 4 and 5. **Your share of the money is**	=	6	11

Seat

Form A

REMAINING SHARE FROM ROUND A

Transfer amount from row 2 of form A1.		1	2,5
Transfer amount from row 2 of form A2.	+	2	3,5
Add amounts from rows 1 and 2. **Remaining share from round A is**	=	3	6

YOUR SHARE FROM ROUND A

Transfer amount from row 4 of form A1.		4	2,5
Transfer amount from row 4 of form A2.	+	5	3,5
Add amounts from rows 4 and 5. **Your share from round A is**	=	6	6

Seat

Form A1

REMAINING SHARE FROM ROUND A **(sliders no. 1-24)**

Money from your task in round A (sliders no. 1-24)		1	5
	x		**50%**
Multiply row 1 with 50%. **Remaining share from round A is (sliders no. 1-24)**	=	2	2,5

YOUR SHARE FROM ROUND A **(sliders no. 1-24)**

Transfer amount from row 1.		3	5
	x		**50%**
Multiply row 3 with 50%. **Your share from round A is (sliders no. 1-24)**	=	4	2.5

Seat

Form A2

REMAINING SHARE FROM ROUND A **(sliders no. 25-48)**

Money from your task in round A (sliders no. 25-48)		1	7
	x		**50%**
Multiply row 1 with 50%. **Remaining share from round A is (sliders no. 25-48)**	=	2	3,5

YOUR SHARE FROM ROUND A **(sliders no. 25-48)**

Transfer amount from row 1.		3	7
	x		**50%**
Multiply row 3 with 50%. **Your share from round A is (sliders no. 25-48)**	=	4	3.5

Seat

Form B

REMAINING SHARE FROM ROUND B

Transfer amount from row 2 of form B1.		1	2
Transfer amount from row 2 of form B2.	+	2	3
Add amounts from rows 1 and 2. **Remaining share from round B is**	=	3	5

YOUR SHARE FROM ROUND B

Transfer amount from row 4 of form B1.		4	2
Transfer amount from row 4 of form B2.	+	5	3
Add amounts from rows 4 and 5. **Your share from round B is**	=	6	5

Seat

Form B1

REMAINING SHARE FROM ROUND B **(sliders no. 1-24)**

Money from your task in round B (sliders no. 1-24)		1	4
	x		**50%**
Multiply row 1 with 50%. **Remaining share from round B is (sliders no. 1-24)**	=	2	2

YOUR SHARE FROM ROUND B **(sliders no. 1-24)**

Transfer amount from row 1.		3	4
	x		**50%**
M Multiply row 3 with 50%. **Your share from round B is (sliders no. 1-24)**	=	4	2

Seat

Form B2

REMAINING SHARE FROM ROUND B **(sliders no. 25-48)**

Money from your task in round B (sliders no. 25-48)		1	6
	x		**50%**
Multiply row 1 with 50%. **Remaining share from round B is (sliders no. 25-48)**	=	2	3

YOUR SHARE FROM ROUND B **(sliders no. 25-48)**

Transfer amount from row 1.		3	6
	x		**50%**
Multiply row 3 with 50%. **Your share from round B is (sliders no. 25-48)**	=	4	3

[*COMPLEX FORM 7/7, 34/34 items*]

D.4 Randomization Checks

Table D.1 shows that the treatment randomization was mostly successful. There exist significant differences with respect to subjects' lab experience. Table D.3 shows the robustness of our main results when controlling for experience.

Controls	Moral/ Simple	Waste/ Simple	Moral/ Complex	Waste/ Complex	F-test
Gender	0.63	0.55	0.65	0.59	0.6008
Age	24.59	24.45	23.00	23.90	0.1112
Study	4.05	3.93	4.25	3.67	0.5316
Math score	2.04	2.18	2.18	2.30	0.5016
Monthly income	3.17	3.59	3.71	3.55	0.6880
Experience	2.31	2.45	2.58	2.65	0.0315
Know	0.03	0.07	0.14	0.21	0.3111

Table D.1 Randomization Checks

Notes: Randomization checks on main control variables. *Study* is a variable that described the field of study. *Math score* is the last high-school grade in math that subjects remembered. *Monthly income* describes a category on monthly available income. *Experience* describes how often a subject has taken part in laboratory experiments. *Know* measures how many of the other participants in the laboratory the subject knows.

	(1) Donation (€)	(2) Donation (%)
Waste	-2.338***	-0.129***
	(0.233)	(0.0178)
Complex	0.455*	0.00677
	(0.236)	(0.0136)
Waste × Complex	-1.339***	-0.0574**
	(0.466)	(0.0267)
Income	0.202***	
	(0.0653)	
Experience	-0.404*	-0.0296**
	(0.210)	(0.0126)
Constant	5.324***	0.537***
	(1.212)	(0.0295)
Observations	296	296
Model	OLS	OLS
R^2	0.251	0.185

Table D.2 Regressions of Compliance Behavior

Notes: Robust standard errors clustered at the subject level in parentheses; $^{*}p < 0.10,^{**} p < 0.05,^{***} p < 0.01$.

D.5 Effort to File the Forms

Nine out of 320 subjects took their forms home such that we cannot assure whether they faced a moral tradeoff when dividing the money. These subjects could have taken home their forms intentionally or mistakenly by putting them in the wrong envelope. These are three subjects in Moral/Simple, two in Moral/ Complex, four in Waste/Simple, and none in Waste/Complex. This behavior is not specific to the treatments (Fisher Exact Test; $p = 0.260$), hence we do not expect treatment specific effort cost functions or any kind of reference points to drive our observed patterns. For the question at hand, overall, subjects donated around 15% of the generated money.

The group of subjects that intentionally left most forms empty and hence did not face the moral tradeoff comprises of two subjects in Moral/Simple, five in Moral/ Complex, and three subjects in both Waste/Simple and Waste/Complex. Again, this pattern is not treatment specific (Fisher Exact Test; $p = 0.756$). These subjects gave on average 20% of the generated money away.

Overall, both groups spent significantly less money to be donated than the overwhelming majority of subjects that filed every single item (MWU; $p < 0.001$). The distribution of shares given is shown in Figure D.4.

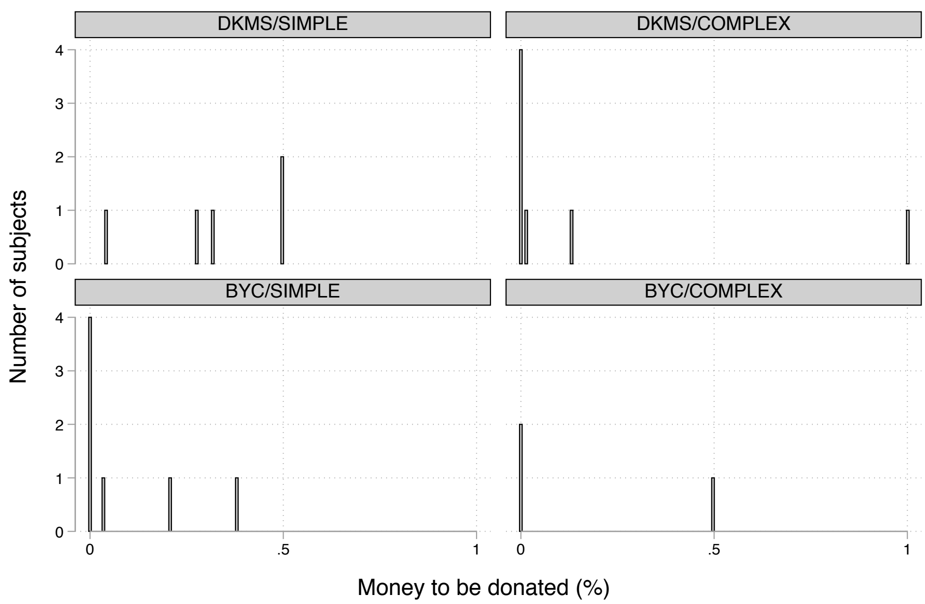

Figure D.4 Donations (in %) and Empty Forms

In our main text we view the effort supply decision to fill the forms as given. One may however argue that the effort supply decision represents a selection into our sample of analysis. In order to address this point, we estimate a simple bivariate selection model (also known as Heckman model, see Heckman (1979)). Here, we re-estimate our linear regression model that presents our main interaction effect (see Table 4.1) by correcting for treatment-specific selection into our sample. Table D.3 shows that the detected effects remain robust. The effect measured using nominal donations slightly increases in level, whereas the respective effect on donation shares decreases slightly.

	(1) Donation (€)	(2) Donation (%)
Waste	-2.221***	-0.137***
	(0.286)	(0.0170)
Complex	0.424	-0.00406
	(0.297)	(0.0125)
Waste × Complex	-1.673***	-0.0454*
	(0.544)	(0.0240)
Income	0.191***	
	(0.0711)	
Constant	4.934***	0.459***
	(1.174)	(0.0105)
Selection on I(sample)		
Waste	0.0149	-0.152
	(0.357)	(0.290)
Complex	-0.259	-0.196
	(0.374)	(0.248)
Waste × Complex	0.807*	0.683*
	(0.432)	(0.380)
Income	0.0460	0.0262
	(0.0287)	(0.0391)
Constant	0.523	1.066
	(0.596)	(0.686)
athrho		
Constant	-1.371***	0.408
	(0.205)	(0.287)
lnsigma		
Constant	1.206***	-1.687***
	(0.0528)	(0.0733)
Observations	318	318
LR-Test	<0.001	0.156
Model	Heckman	Heckman

Table D.3 Regressions of Compliance Behavior and Selection Effects

Notes: Robustness of compliance behavior on the intensive margin. Notice that two subjects that did not position a single slider correctly are excluded as the slider task was performed before knowing about treatment-specific information. Clustered standard errors (session level) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.6 Identification

In the following, we state our hypotheses where the variables x_A, x_B, x_C, x_D (all ≥ 0) represent the average amount of money to be donated to the recipient party in the respective treatment groups (A: MORAL and SIMPLE, B: MORAL and COMPLEX, C: WASTE and SIMPLE, D: WASTE and COMPLEX).

textbfDirect complexity effect
Subjects devote less money to be donated to the other party when the forms are complex rather than simple, such that $(x_A + x_C) - (x_B + x_D) > 0$.

Appropriation effect
Subjects devote less money to be donated to the BYC compared to the DKMS, such that $(x_A + x_B) - (x_C + x_D) > 0$.

Employing more complex forms induces incentives for non-compliance due to higher cognitive effort costs which should lead to less compliance.[1] We lower the moral costs of non-compliance with donating money to the BYC as compared to the DKMS which is also expected to have a negative effect on compliance.

Our hypotheses can be conceptualized within the following simple framework. A decision maker is assumed to maximize his utility U over the share of money to be donated to the other party $x \in [0, 1]$.

$$\max_{\arg x} U = (1 - x) + F[C(i), m] \cdot (-x) - t \cdot G(x - g). \tag{D.1}$$

The utility function consists of three parts. First, $(1 - x)$ is the (consumption) utility from the money kept. Second, $F[C(\cdot), m]$ describes the context effects on the agents decision to optimally choose x. Here, $m \in [0, 1]$ is the subject's social concerns and $C(\cdot)$ represents complexity as a function of the number of different items to file (i). Thirdly, $G(\cdot)$ represents the duty to comply to the rule g with intensity factor t. The first and

[1] Complexity is likely to be a hybrid of decision time, cognitive effort, and depletion that all affect compliance in the same direction.

second order conditions are, respectively:

$$FOC: \qquad \frac{\partial U}{\partial x} = -F[C(\cdot), m] - t\frac{\partial G(\cdot)}{\partial x} = 0 \tag{D.2}$$

$$SOC: \qquad \frac{\partial^2 U}{\partial x^2} = -t\frac{\partial^2 G(\cdot)}{\partial x^2} < 0. \tag{D.3}$$

We easily see that there is a tradeoff between our context variables and the duty to comply to the rule. Using the implicit function theorem we can show that our Hypotheses **??** and **??** can be expressed by $\frac{dx}{dC(\cdot)} = \frac{\frac{\partial F(\cdot)}{\partial C(\cdot)}}{-t\frac{\partial^2 G(\cdot)}{\partial x^2}} < 0$ and $\frac{dx}{dm} = \frac{\frac{\partial F(\cdot)}{\partial m}}{-t\frac{\partial^2 G(\cdot)}{\partial x^2}} < 0$, respectively.

However, the exact treatment response depends on the functional form of $F[C(\cdot), m]$ that we aim to better understand with this paper. We cautiously formulate the following alternative hypothesis on $F[C(\cdot), m]$ when we conjecture that $\frac{\partial^2 F}{\partial C \partial m} < 0$. I.e., the negative interaction effect of complexity and appropriation gets weaker when moral costs increase. We propose that the behavioral mechanism underlying this relation is that complexity can reduce moral costs from non-compliance. In our simple decision framework, this is reflected by a negative complexity effect that diminishes with moral costs. An alternative interpretation would be that the base motivation of employees is reduced by the appropriation of donations such that complexity effects loom larger when the money is used for a wasteful purpose.

Interaction effect
The effects of complexity depend on the recipient (i.e., the moral context). In particular, the negative effects of complexity are stronger for the BYC than for the DKMS, such that $[(x_A - x_B) - (x_C - x_D)] < 0$.

Bibliography

Abeler, Johannes and Simon Jaeger, "Complex tax incentives," *American Economic Journal: Economic Policy*, 2015, *7* (3), pp. 1–28.

— , Daniele Nosenzo, and Collin Raymond, "Preferences for truth-telling," *Econometrica*, 2019, *87* (4), pp. 1115–1153.

Agranov, Marina and Andrew Schotter, "Ignorance is bliss: an experimental study of the use of ambiguity and vagueness in the coordination games with asymmetric payoffs," *American Economic Journal: Microeconomics*, 2012, *4* (2), pp. 77–103.

Algan, Yann, Yochai Benkler, Emeric Henry, and Jérôme Hergueux, "Social motives and the organization of production: Experimental evidence from open source software," *Working Paper*, 2014.

— , — , Mayo Fuster Morell, and Jérôme Hergueux, "Cooperation in a peer production economy experimental evidence from wikipedia," *Working Paper*, 2013.

Alpizar, Francisco, Fredrik Carlsson, and Olof Johansson-Stenman, "Anonymity, reciprocity, and conformity: Evidence from voluntary contributions to a national park in Costa Rica," *Journal of Public Economics*, 2008, *92* (5-6), pp. 1047–1060.

Amar, Akhil Reed and Renee Lettow, "Fifth amendment first principles: The self-incrimination clause," *Michigan Law Review*, 1995, *93* (5), pp. 857–928.

Ambrus, Attila and Parag Pathak, "Cooperation over finite horizons: A theory and experiments," *Journal of Public Economics*, 2011, *95* (7-8), pp. 500–512.

Anderson, Lisa, Jennifer Mellor, and Jeffrey Milyo, "Social capital and contributions in a public-goods experiment," *American Economic Review*, 2004, *94* (2), pp. 373–376.

Ashforth, Blake and Fred Mael, "Social identity theory and the organization," *Academy of Management Review*, 1989, *14* (1), pp. 20–39.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul, "Cooperation in collective action," *Economics of Transition*, 2005, *13* (3), pp. 473–498.

_ , _ , and _ , "Field experiments with firms," *Journal of Economic Perspectives*, 2011, *25* (3), pp. 63–82.

_ , _ , and _ , "Team incentives: Evidence from a firm level experiment," *Journal of the European Economic Association*, 2013, *11* (5), pp. 1079–1114.

Barr, Abigail and Pieter Serneels, "Reciprocity in the workplace," *Experimental Economics*, 2009, *12* (1), pp. 99–112.

Barr, Abigal, Jean Ensminger, and Jeffrey Johnson, "Social networks and trust in cross-cultural economic experiments," in "Whom can we trust?," New York: Russell Sage Foundation, 2009.

Beersma, Bianca, John Hollenbeck, Stephen Humphrey, Henry Moon, Donald Conlon, and Daniel Ilgen, "Cooperation, competition, and team performance: Toward a contingency approach," *Academy of Management Journal*, 2003, *46* (5), pp. 572–590.

Bénabou, Roland and Jean Tirole, "Intrinsic and extrinsic motivation," *Review of Economic Studies*, 2003, *70* (3), pp. 489–520.

Bénabou, Roland and Jean Tirole, "Laws and norms," *NBER Working Paper No. 17579*, 2011.

_ , Armin Falk, and Jean Tirole, "Narratives, imperatives and moral reasoning," *NBER Working Paper No. 24798*, 2018.

Benndorf, Volker, Dorothea Kübler, and Hans-Theo Normann, "Privacy concerns, voluntary disclosure of information, and unraveling: An experiment," *European Economic Review*, 2015, *75*, pp. 43–59.

Benz, Matthias and Stephan Meier, "Do people behave in experiments as in the field? Evidence from donations," *Experimental Economics*, 2008, *11* (3), pp. 268–281.

Benzarti, Youssef, "How taxing is tax filing? Using revealed preferences to estimate compliance costs," *American Economic Journal: Economic Policy*, 2020, *forthcomming.*

Bernheim, Douglas and Dmitry Taubinsky, "Behavioral public economics," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 1, Elsevier, 2018, pp. 381–516.

Bieberstein, Frauke Von, Jonas Gehrlein, and Anna Güntner, "Teamwork revisited: Social preferences and knowledge acquisition in the field," *Journal of Business Economics*, 2020, pp. 1–24.

Blumenthal, Marsha and Joel Slemrod, "The compliance cost of the US individual income tax system: A second look after tax reform," *National Tax Journal*, 1992, *45* (2), pp. 185–202.

Board, Oliver, Andreas Blume, and Kohei Kawamura, "Noisy talk," *Theoretical Economics*, December 2007, *2* (4), pp. 395–440.

Bolton, Gary and Axel Ockenfels, "ERC: A theory of equity, reciprocity, and competition," *American Economic Review*, 2000, *90* (1), pp. 166–193.

Bowles, Samuel, "Endogenous preferences: The cultural consequences of markets and other economic institutions," *Journal of Economic Literature*, 1998, *36* (1), pp. 75–111.

__ **and Sandra Polanía-Reyes**, "Economic incentives and social preferences: substitutes or complements?," *Journal of Economic Literature*, 2012, *50* (2), pp. 368–425.

__ **, Herbert Gintis, and Melissa Osborne**, "Incentive-enhancing preferences: Personality, behavior, and earnings," *American Economic Review*, 2001, *91* (2), pp. 155–158.

Bremzen, Andrei, Elena Khokhlova, Anton Suvorov, and Jeroen Van de Ven, "Bad news: An experimental study on the informational effects of rewards," *Review of Economics and Statistics*, 2015, *97* (1), pp. 55–70.

Brodeur, Abel, Mathias Lé, Marc Sangnier, and Yanos Zylberberg, "Star wars: The empirics strike back," *American Economic Journal: Applied Economics*, 2016, *8* (1), pp. 1–32.

Brown, Alexander and Daniel Fragiadakis, "Can individuals correctly respond to strategic attribute disclosure?," *Working Paper*, 2018.

Brown, Alexander L, Colin F Camerer, and Dan Lovallo, "To review or not to review? Limited strategic thinking at the movie box office," *American Economic Journal: Microeconomics*, 2012, *4* (2), pp. 1–26.

__ **,** __ **, and** __ , "Estimating structural models of equilibrium and cognitive hierarchy thinking in the field: The case of withheld movie critic reviews," *Management Science*, 2013, *59* (3), pp. 733–747.

Brown, Jennifer, Tanjim Hossain, and John Morgan, "Shrouded attributes and information suppression: Evidence from the field," *Quarterly Journal of Economics*, 2010, *125* (2), pp. 859–876.

Burks, Stephen and Erin Krupka, "A multimethod approach to identifying norms and normative expectations within a corporate hierarchy: Evidence from the financial services industry," *Management Science*, 2012, *58* (1), pp. 203–217.

_ , Daniele Nosenzo, Jon Anderson, Matthew Bombyk, Derek Ganzhorn, Lorenz Götte, and Aldo Rustichini, "Lab measures of other-regarding preferences can predict some related on-the-job behavior: Evidence from a large scale field experiment," *Working Paper*, 2016.

_ , Jeffrey Carpenter, and Lorenz Goette, "Performance pay and worker cooperation: Evidence from an artefactual field experiment," *Journal of Economic Behavior & Organization*, 2009, *70* (3), pp. 458–469.

Cain, Daylian, George Loewenstein, and Don Moore, "The dirt on coming clean: Perverse effects of disclosing conflicts of interest," *Journal of Legal Studies*, 2005, *34* (1), pp. 1–25.

Camerer, Colin, "Progress in behavioral game theory," *Journal of Economic Perspectives*, 1997, *11* (4), pp. 167–188.

_ and Ulrike Malmendier, "Behavioral organizational economics," in "Behavioral Economics and Its Applications," Citeseer, 2007, pp. 235–290.

Cardenas, Juan-Camilo, "Real wealth and experimental cooperation: Experiments in the field lab," *Journal of Development Economics*, 2003, *70* (2), pp. 263–289.

Cardenas, Juan Camilo and Jeffrey Carpenter, "Three themes on field experiments and economic development," in "Field Experiments in Economics," JAI Press, Greenwich, 2005, pp. 71–124.

Cardinaels, Eddy and Huaxiang Yin, "Think twice before going for incentives: Social norms and the principal's decision on compensation contracts," *Journal of Accounting Research*, 2015, *53* (5), pp. 985–1015.

Carpenter, Jeffrey and Erika Seki, "Do social preferences increase productivity? Field experimental evidence from fishermen in Toyama Bay," *Economic Inquiry*, 2011, *49* (2), pp. 612–630.

Charness, Gary and Marie-Claire Villeval, "Cooperation and competition in intergenerational experiments in the field and the laboratory," *American Economic Review*, 2009, *99* (3), pp. 956–978.

_ **and Matthew Rabin**, "Understanding social preferences with simple tests," *Quarterly Journal of Economics*, 2002, *117* (3), pp. 817–869.

Chaudhuri, Ananish, "Sustaining cooperation in laboratory public goods experiments: A selective survey of the literature," *Experimental Economics*, 2011, *14* (1), pp. 47–83.

Chen, Ying, "Perturbed communication games with honest senders and naive receivers," *Journal of Economic Theory*, 2011, *146* (2), pp. 401–424.

Chetty, Raj, "Behavioral economics and public policy: A pragmatic perspective," *American Economic Review*, 2015, *105* (5), pp. 1–33.

Cohn, Alain, Ernst Fehr, and Michel André Maréchal, "Business culture and dishonesty in the banking industry," *Nature*, 2014, *516* (7529), pp. 86–89.

Crawford, Vincent and Joel Sobel, "Strategic information transmission," *Econometrica*, 1982, pp. 1431–1451.

Croson, Rachel and Jen Yue Shang, "The impact of downward social information on contribution decisions," *Experimental Economics*, 2008, *11* (3), pp. 221–233.

Czermak, Simon, Francesco Feri, Daniela Glätzle-Rützler, and Matthias Sutter, "How strategic are children and adolescents? Experimental evidence from normal-form games," *Journal of Economic Behavior & Organization*, 2016, *128*, pp. 265–285.

Danilov, Anastasia and Dirk Sliwka, "Can contracts signal social norms? Experimental evidence," *Management Science*, 2017, *63* (2), pp. 459–476.

DellaVigna, Stefano, "Psychology and economics: Evidence from the field," *Journal of Economic literature*, 2009, *47* (2), pp. 315–72.

der Weele, Joël Van, "The signaling power of sanctions in social dilemmas," *Journal of Law, Economics, & Organization*, 2012, *28* (1), pp. 103–126.

Dertwinkel-Kalt, Markus, Mats Köster, and Matthias Sutter, "To buy or not to buy? Shrouding and partitioning of prices in an online shopping field experiment," *CESifo Working Paper No. 7475*, 2019.

Deserranno, Erika, "Financial incentives as signals: experimental evidence from the recruitment of village promoters in Uganda," *American Economic Journal: Applied Economics*, 2019, *11* (1), pp. 277–317.

Deversi, Marvin, Martin Kocher, and Christiane Schwieren, "Cooperation in a company: A large-scale experiment," *Working Paper*, 2020.

_, _, Robert Schmidt, Christoph Schweighofer, and Christiane Schwieren, "Norms of cooperation in a company," *Working Paper*, 2020.

Dickhaut, John, Margaret Ledyard, Arijit Mukherji, and Haresh Sapra, "Information management and valuation: an experimental investigation," *Games and Economic Behavior*, 2003, *44* (1), pp. 26–53.

Dirks, Kurt, "The effects of interpersonal trust on work group performance," *Journal of Applied Psychology*, 1999, *84* (3), pp. 445–455.

_ and Donald Ferrin, "The role of trust in organizational settings," *Organization Science*, 2001, *12* (4), pp. 450–467.

Dohmen, Thomas and Armin Falk, "Performance pay and multidimensional sorting: Productivity, preferences, and gender," *American Economic Review*, 2011, *101* (2), pp. 556–90.

Dranove, David and Ginger Zhe Jin, "Quality disclosure and certification: Theory and practice," *Journal of Economic Literature*, 2010, *48* (4), pp. 935–963.

Dwenger, Nadja, Henrik Kleven, Imran Rasul, and Johannes Rincke, "Extrinsic and intrinsic motivations for tax compliance: Evidence from a field experiment in Germany," *American Economic Journal: Economic Policy*, 2016, *8* (3), pp. 203–232.

Ellison, Glenn and Sara Fisher Ellison, "Search, obfuscation, and price elasticities on the internet," *Econometrica*, 2009, *77* (2), pp. 427–452.

Englmaier, Florian and Georg Gebhardt, "Social dilemmas in the laboratory and in the field," *Journal of Economic Behavior & Organization*, 2016, *128*, pp. 85–96.

Essl, Andrea, Michael Kosfeld, Markus Kröll, and Frauke Von Bieberstein, "Sales performance and social preferences," *Working Paper*, 2018.

Exley, Christine and Judd Kessler, "Motivated Errors," *NBER Working Paper No. 26595*, 2019.

Exley, Christine L, "Using charity performance metrics as an excuse not to give," *Management Science*, 2020, *66* (2), pp. 553–563.

Eyster, Erik and Matthew Rabin, "Cursed equilibrium," *Econometrica*, 2005, *73* (5), pp. 1623–1672.

Falk, Armin and James Heckman, "Lab experiments are a major source of knowledge in the social sciences," *Science*, 2009, *326* (5952), pp. 535–538.

_, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, "Global evidence on economic preferences," *Quarterly Journal of Economics*, 2018, *133* (4), pp. 1645–1692.

Fehr, Ernst, "Behavioral foundations of corporate culture," *UBS Center Public Paper No. 7*, 2018.

_ and Andreas Leibbrandt, "A field study on cooperativeness and impatience in the tragedy of the commons," *Journal of Public Economics*, 2011, *95* (9-10), pp. 1144–1155.

_ and Bettina Rockenbach, "Detrimental effects of sanctions on human altruism," *Nature*, 2003, *422* (6928), pp. 137–140.

_ and Ivo Schurtenberger, "Normative foundations of human cooperation," *Nature Human Behaviour*, 2018, *2* (7), pp. 458–468.

_ and John List, "The hidden costs and returns of incentives – trust and trustworthiness among CEOs," *Journal of the European Economic Association*, 2004, *2* (5), pp. 743–771.

_ and Klaus Schmidt, "A theory of fairness, competition, and cooperation," *Quarterly Journal of Economics*, 1999, *114* (3), pp. 817–868.

_ and Urs Fischbacher, "Why social preferences matter–the impact of non-selfish motives on competition, cooperation and incentives," *Economic Journal*, 2002, *112* (478), pp. 1–33.

Feldman, Naomi E, Peter Katuščák, and Laura Kawano, "Taxpayer confusion: Evidence from the child tax credit," *American Economic Review*, 2016, *106* (3), pp. 807–835.

Ferman, Bruno, "Reading the fine print: information disclosure in the Brazilian credit card market," *Management Science*, 2015, *62* (12), pp. 3534–3548.

Fischbacher, Urs, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 2007, *10* (2), pp. 171–178.

__ **and Simon Gächter**, "Social preferences, beliefs, and the dynamics of free riding in public goods experiments," *American Economic Review*, 2010, *100* (1), pp. 541–56.

__ **, Simon Gächter, and Ernst Fehr**, "Are people conditionally cooperative? Evidence from a public goods experiment," *Economics Letters*, 2001, *71* (3), pp. 397–404.

Forsythe, Robert, Mark Isaac, and Thomas Palfrey, "Theories and tests of "blind bidding" in sealed-bid auctions," *RAND Journal of Economics*, 1989, pp. 214–238.

Frey, Bruno and Stephan Meier, "Social comparisons and pro-social behavior: Testing" conditional cooperation" in a field experiment," *American Economic Review*, 2004, *94* (5), pp. 1717–1722.

Friebel, Guido and Wendelin Schnedler, "Team governance: Empowerment or hierarchical control," *Journal of Economic Behavior & Organization*, 2011, *78* (1-2), pp. 1–13.

Gabaix, Xavier and David Laibson, "Shrouded attributes, consumer myopia, and information suppression in competitive markets," *Quarterly Journal of Economics*, 2006, *121* (2), pp. 505–540.

Gächter, Simon, "Conditional cooperation: Behavioral regularities from the lab and the field and their policy implications," in "Psychology and Economics: A Promising New Cross-Disciplinary Field," MIT Press, Cambridge, 2007, pp. 19–50.

Gächter, Simon and Elke Renner, "The effects of (incentivized) belief elicitation in public goods experiments," *Experimental Economics*, 2010, *13* (3), pp. 364–377.

__ **, Felix Kölle, and Simone Quercia**, "Reciprocity and the tragedies of maintaining and providing the commons," *Nature Human Behaviour*, 2017, *1* (9), pp. 650–656.

Galbiati, Roberto, Karl Schlag, and Joël Van Der Weele, "Sanctions that signal: An experiment," *Journal of Economic Behavior & Organization*, 2013, *94*, pp. 34–51.

Galizzi, Matteo and Daniel Navarro-Martínez, "On the external validity of social preference games: A systematic lab-field study," *Management Science*, 2019, *65* (3), pp. 976–1002.

Gill, David and Victoria Prowse, "Measuring costly effort using the slider task," *Journal of Behavioral and Experimental Finance*, 2019, *21*, pp. 1–9.

Gino, Francesca, Maurice Schweitzer, Nicole Mead, and Dan Ariely, "Unable to resist temptation: How self-control depletion promotes unethical behavior," *Organizational Behavior and Human Decision Processes*, 2011, *115* (2), pp. 191–203.

Gittell, Jody Hoffer, "Organizing work to support relational co-ordination," *International Journal of Human Resource Management*, 2000, *11* (3), pp. 517–539.

Gneezy, Uri and Alex Imas, "Lab in the field: Measuring preferences in the wild," in "Handbook of Economic Field Experiments," Vol. 1, Elsevier, 2017, pp. 439–464.

_ , Andreas Leibbrandt, and John A List, "Ode to the sea: Workplace organizations and norms of cooperation," *Economic Journal*, 2016, *126* (595), pp. 1856–1883.

_ , Elizabeth Keenan, and Ayelet Gneezy, "Avoiding overhead aversion in charity," *Science*, 2014, *346* (6209), pp. 632–635.

_ , Stephan Meier, and Pedro Rey-Biel, "When and why incentives (don't) work to modify behavior," *Journal of Economic Perspectives*, 2011, *25* (4), pp. 191–210.

Government Accountability Office, "Tax gap: Complexity and taxpayer compliance," *Policy Report*, 2017.

Grant, Adam M, *Give and take: A revolutionary approach to success*, Penguin, 2013.

Gratton, Lynda, "How to foster a cooperative culture," *Harvard Business Review*, 2009.

_ , *The shift: The future of work is already here*, HarperCollins UK, 2011.

Greiner, Ben, "Subject pool recruitment procedures: Organizing experiments with ORSEE," *Journal of the Economic Science Association*, 2015, *1* (1), pp. 114–125.

Grossman, Sanford J, "The informational role of warranties and private disclosure about product quality," *Journal of Law and Economics*, 1981, *24* (3), pp. 461–483.

_ and Oliver D Hart, "Disclosure laws and takeover bids," *Journal of Finance*, 1980, *35* (2), pp. 323–334.

Hagenbach, Jeanne and Eduardo Perez-Richet, "Communication with Evidence in the Lab," *Games and Economic Behavior*, 2018, *112*, pp. 139–165.

_ and Frédéric Koessler, "Simple versus rich language in disclosure games," *Review of Economic Design*, 2017, pp. 1–13.

Haisley, Emily and Roberto Weber, "Self-serving interpretations of ambiguity in other-regarding behavior," *Games and Economic Behavior*, 2010, *68* (2), pp. 614–625.

Hallsworth, Michael, John A List, Robert D Metcalfe, and Ivo Vlaev, "The behavioralist as tax collector: Using natural field experiments to enhance tax compliance," *Journal of Public Economics*, 2017, *148*, pp. 14–31.

Hamilton, Barton, Jack Nickerson, and Hideo Owan, "Team incentives and worker heterogeneity: An empirical analysis of the impact of teams on productivity and participation," *Journal of Political Economy*, 2003, *111* (3), pp. 465–497.

Harrison, Glenn and John List, "Field experiments," *Journal of Economic Literature*, 2004, *42* (4), pp. 1009–1055.

Harrison, Glenn W, Jimmy Martínez-Correa, and J Todd Swarthout, "Eliciting subjective probabilities with binary lotteries," *Journal of Economic Behavior & Organization*, 2014, *101*, pp. 128–140.

Hartigan, John and PM Hartigan, "The dip test of unimodality," *Annals of Statistics*, 1985, pp. 70–84.

Heidhues, Paul and Botond Kőszegi, "Behavioral industrial organization," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 1, Elsevier, 2018, pp. 517–612.

Hettich, Walter and Stanley Winer, "Economic and political foundations of tax structure," *American Economic Review*, 1988, pp. 701–712.

Holmstrom, Bengt, "Moral hazard in teams," *Bell Journal of Economics*, 1982, pp. 324–340.

Hossain, Tanjim and John A List, "The behavioralist visits the factory: Increasing productivity using simple framing manipulations," *Management Science*, 2012, *58* (12), pp. 2151–2167.

_ **and Ryo Okui**, "The binarized scoring rule," *Review of Economic Studies*, 2013, *80* (3), pp. 984–1001.

Houser, Daniel, Erte Xiao, Kevin McCabe, and Vernon Smith, "When punishment fails: Research on sanctions, intentions and non-cooperation," *Games and Economic Behavior*, 2008, *62* (2), pp. 509–532.

Irlenbusch, Bernd and Gabriele Ruchala, "Relative rewards within team-based compensation," *Labour Economics*, 2008, *15* (2), pp. 141–167.

Ispano, Alessandro and Peter Schwardmann, "Competition over Cursed Consumers," *CESifo Working Paper No. 7046*, 2018.

Itoh, Hideshi, "Incentives to help in multi-agent situations," *Econometrica*, 1991, pp. 611–636.

Jin, Ginger Zhe, "Competition and disclosure incentives: an empirical study of HMOs," *RAND Journal of Economics*, 2005, pp. 93–112.

__ **, Michael Luca, and Daniel Martin**, "Is no news (perceived as) bad news? An experimental investigation of information disclosure," *NBER Working Paper No. 21099*, 2018.

__ **,** __ **, and** __ , "Complex disclosure," *Working Paper*, 2019.

Kahneman, Daniel and Amos Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica*, 1979, *47* (2), pp. 263–291.

Kandel, Eugene and Edward Lazear, "Peer pressure and partnerships," *Journal of Political Economy*, 1992, *100* (4), pp. 801–817.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani, "Credulity, lies, and costly talk," *Journal of Economic Theory*, May 2007, *134* (1), pp. 93–116.

King, Ronald and David Wallin, "Voluntary disclosures when seller's level of information is unknown," *Journal of Accounting Research*, 1991, pp. 96–108.

Kocher, Martin, Peter Martinsson, Dominik Matzat, and Conny Wollbrant, "The role of beliefs, trust, and risk in contributions to a public good," *Journal of Economic Psychology*, 2015, *51*, pp. 236–244.

__ **, Todd Cherry, Stephan Kroll, Robert Netzer, and Matthias Sutter**, "Conditional cooperation on three continents," *Economics Letters*, 2008, *101* (3), pp. 175–178.

Konow, James, "Fair shares: Accountability and cognitive dissonance in allocation decisions," *American Economic Review*, 2000, *90* (4), pp. 1072–1091.

Kosfeld, Michael and Ferdinand Von Siemens, "Competition, cooperation, and corporate culture," *RAND Journal of Economics*, 2011, *42* (1), pp. 23–43.

Kőszegi, Botond and Matthew Rabin, "A model of reference-dependent preferences," *Quarterly Journal of Economics*, 2006, *121* (4), pp. 1133–1165.

Krawczyk, Micha l, "The search for significance: a few peculiarities in the distribution of P-values in experimental psychology literature," *PloS One*, 2015, *10* (6), e0127872.

Laibson, David, "Golden eggs and hyperbolic discounting," *Quarterly Journal of Economics*, 1997, *112* (2), pp. 443–478.

Ledyard, John, "Public goods: A survey of experimental research," *Handbook of Experimental Economics*, 1995, *Vol. 1.*

Leib, Margarita, Andrea Pittarello, Tom Gordon-Hecker, Shaul Shalvi, and Marieke Roskes, "Loss framing increases self-serving mistakes (but does not alter attention)," *Journal of Experimental Social Psychology*, 2019, *85*, pp. 103880.

Leibbrandt, Andreas, "Are social preferences related to market performance?," *Experimental Economics*, 2012, *15* (4), pp. 589–603.

Levitt, Steven and John List, "What do laboratory experiments measuring social preferences reveal about the real world?," *Journal of Economic Perspectives*, 2007, *21* (2), pp. 153–174.

Lewis, Gregory, "Asymmetric information, adverse selection and online disclosure: The case of eBay motors," *American Economic Review*, 2011, *101* (4), pp. 1535–1546.

Li, Ying Xue and Burkhard C Schipper, "Strategic reasoning in persuasion games: An Experiment," *Games and Economic Behavior*, 2020, *forthcomming*.

List, John A and David Lucking-Reiley, "The effects of seed money and refunds on charitable giving: Experimental evidence from a university capital campaign," *Journal of Political Economy*, 2002, *110* (1), pp. 215–233.

Loewenstein, George, Cass Sunstein, and Russell Golman, "Disclosure: Psychology changes everything," *Annual Review of Economics*, 2014, *6* (1), pp. 391–419.

Luca, Michael and Jonathan Smith, "Strategic disclosure: The case of business school rankings," *Journal of Economic Behavior & Organization*, 2015, *112*, pp. 17–25.

Luttmer, Erzo and Monica Singhal, "Tax morale," *Journal of Economic Perspectives*, 2014, *28* (4), pp. 149–168.

Mathios, Alan, "The impact of mandatory disclosure laws on product choices: An analysis of the salad dressing market," *Journal of Law and Economics*, 2000, *43* (2), pp. 651–678.

Matthews, Steven and Andrew Postlewaite, "Quality testing and disclosure," *RAND Journal of Economics*, Autumn 1985, *16* (3), pp. 328–340.

Mead, Nicole, Roy Baumeister, Francesca Gino, Maurice Schweitzer, and Dan Ariely, "Too tired to tell the truth: Self-control resource depletion and dishonesty," *Journal of Experimental Social Psychology*, 2009, *45* (3), pp. 594–597.

Milgrom, Paul, "Good news and bad news: Representation theorems and applications," *Bell Journal of Economics*, 1981, pp. 380–391.

Milgrom, Paul R and John Roberts, "Relying on the information of interested parties," *RAND Journal of Economics*, Spring 1986, *17* (1), pp. 18–32.

Newby, Jennifer and Rupert Klein, "Competitiveness reconceptualized: Psychometric development of the competitiveness orientation measure as a unified measure of trait competitiveness," *The Psychological Record*, 2014, *64* (4), pp. 879–895.

Nikiforakis, Nikos, "Punishment and counter-punishment in public good games: Can we really govern ourselves?," *Journal of Public Economics*, 2008, *92* (1-2), pp. 91–112.

Ostrom, Elinor, *Governing the commons: The evolution of institutions for collective action*, Cambridge University Press, 1990.

Ottaviani, Marco and Francesco Squintani, "Naive audience and communication bias," *International Journal of Game Theory*, December 2006, *35* (1), pp. 129–150.

Polinsky, Mitchell and Steven Shavell, "Mandatory versus voluntary disclosure of product risks," *Journal of Law, Economics, and Organization*, 2010, pp. 360–379.

Rabin, Matthew, "A perspective on psychology and economics," *European Economic Review*, 2002, *46* (4-5), pp. 657–685.

Rammstedt, Beatrice, Christoph Kemper, Mira Céline Klein, Constanze Beierlein, and Anastassiya Kovaleva, "Eine kurze Skala zur Messung der fünf Dimensionen der Persönlichkeit: Big-Five-Inventory-10 (BFI-10)," *Methoden, Daten, Analysen*, 2013, *7* (2), pp. 233–249.

Rosen, Sherwin, "The theory of equalizing differences," *Handbook of Labor Economics*, 1986, *1*, pp. 641–692.

Roth, Alvin, "Laboratory experimentation in economics," *Economics & Philosophy*, 1986, *2* (2), pp. 245–273.

_ **and Michael Malouf**, "Game-theoretic models and the role of information in bargaining," *Psychological Review*, 1979, *86* (6), pp. 574–594.

Ru, Hong and Antoinette Schoar, "Do credit card companies screen for behavioral biases?," *NBER Working Paper No. 22360*, 2016.

Ruff, Christian and Ernst Fehr, "The neurobiology of rewards and values in social decision making," *Nature Reviews Neuroscience*, 2014, *15* (8), pp. 549–562.

Rustagi, Devesh, Stefanie Engel, and Michael Kosfeld, "Conditional cooperation and costly monitoring explain success in forest commons management," *Science*, 2010, *330* (6006), pp. 961–965.

Schnedler, Wendelin and Radovan Vadovic, "Legitimacy of control," *Journal of Economics & Management Strategy*, 2011, *20* (4), pp. 985–1009.

Schulz, Peter and Wolff Schlotz, "Trierer Inventar zur Erfassung von chronischem Sre (TICS): Skalenkonstruktion, teststatistische Überprüfung und Validierung der Skala Arbeitsüberlastung.," *Diagnostica*, 1999, pp. 8–19.

Serra-Garcia, Marta, Eric van Damme, and Jan Potters, "Hiding an inconvenient truth: Lies and vagueness," *Games and Economic Behavior*, 2011, *73* (1), pp. 244–261.

Simon, Herbert, "A behavioral model of rational choice," *Quarterly Journal of Economics*, 1955, *69* (1), pp. 99–118.

Slemrod, Joel and Nikki Sorum, "The compliance cost of the US individual income tax system," *NBER Working Paper No. 1401*, 1984.

Sliwka, Dirk, "Trust as a signal of a social norm and the hidden costs of incentive schemes," *American Economic Review*, 2007, *97* (3), pp. 999–1012.

Spiegler, Ran, "Choice complexity and market competition," *Annual Review of Economics*, 2016, *8*, pp. 1–25.

Stanford Law Review, "Waiver of the privilege against self incrimination," 1962, *14* (4), pp. 811–826.

Stoop, Jan, Charles Noussair, and Daan Van Soest, "From the lab to the field: Cooperation among fishermen," *Journal of Political Economy*, 2012, *120* (6), pp. 1027–1056.

Swift, Samuel, Don Moore, Zachariah Sharek, and Francesca Gino, "Inflated applicants: Attribution errors in performance evaluation by professionals," *PLoS One*, 2013, *8* (7), e69258.

Taubinsky, Dmitry and Alex Rees-Jones, "Attention variation and welfare: theory and evidence from a tax salience experiment," *Review of Economic Studies*, 2017, *85* (4), pp. 2462–2496.

Thaler, Richard, "Toward a positive theory of consumer choice," *Journal of Economic Behavior & Organization*, 1980, *1* (1), pp. 39–60.

Thaler, Richard H and Shlomo Benartzi, "Save more tomorrow: Using behavioral economics to increase employee saving," *Journal of Political Economy*, 2004, *112* (S1), pp. 164–187.

Torgler, Benno, "Tax morale, rule-governed behaviour and trust," *Constitutional Political Economy*, 2003, *14* (2), pp. 119–140.

Vaillancourt, François, Marylène Roy, and Charles Lammam, "Measuring tax complexity in Canada," *Fraser Research Bulletin*, 2015, pp. 1–10.

Voors, Maarten, Erwin Bulte, Andreas Kontoleon, John List, and Ty Turley, "Using artefactual field experiments to learn about the incentives for sustainable forest use in developing economies," *American Economic Review*, 2011, *101* (3), pp. 329–33.

_ **, Ty Turley, Andreas Kontoleon, Erwin Bulte, and John List**, "Exploring whether behavior in context-free experiments is predictive of behavior in the field: Evidence from lab and field experiments in rural Sierra Leone," *Economics Letters*, 2012, *114* (3), pp. 308–311.

Wood, Daniel H, "Communication-enhancing vagueness," *Working Paper*, 2016.

Yale Law Journal, "The privilege against self-incrimination: The doctrine of waiver," 1952, *61* (1), pp. 105–110.

Eidesstattliche Versicherung

Ich versichere hiermit eidesstattlich, dass ich die vorliegende Arbeit selbstständig und ohne fremde Hilfe verfasst habe. Die aus fremden Quellen direkt oder indirekt übernommenen Gedanken sowie mir gegebene Anregungen sind als solche kenntlich gemacht. Die Arbeit wurde bisher keiner anderen Prüfungsbehörde vorgelegt und auch noch nicht veröffentlicht. Sofern ein Teil der Arbeit aus bereits veröffentlichten Papers besteht, habe ich dies ausdrücklich angegeben.

Datum:

Unterschrift: