
MARKET DESIGN - ESSAYS ON THE OPTIMAL STRUCTURE OF
AUCTIONS AND MATCHING MARKETS

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

2023

vorgelegt von

Christoph Robert Schwaiger

Referent:

Prof. Dr. Klaus M. Schmidt

Koreferent:

Prof. Dr. Matthias Lang

Promotionsabschlussberatung:

01. Februar 2023

Datum der mündlichen Prüfung: 27. Januar 2023

Namen der Berichterstatter: Klaus Schmidt, Matthias Lang, Florian Englmaier

Acknowledgements

This dissertation is the result of continuous support and constant encouragement from family, friends, colleagues, and superiors.

I am very grateful to Klaus Schmidt for his excellent support. He is always approachable and interested in discussing research. His comments and remarks have greatly improved each of my research projects. I am also very grateful to my second supervisor, Matthias Lang, who always supported me in every way and from whom I was able to learn a lot. I would also like to thank Florian Englmaier, who agreed to serve as my third supervisor. His ability to bring people together has greatly enriched both my time and that of many colleagues.

I must especially thank my co-author Timm Opitz. Working on both projects with him was always fun, and his ideas and diligence significantly improved this dissertation. Throughout the Ph.D., I was lucky to work with great and humorous colleagues at the chair of economic theory. It was a pleasure to work together with Benjamin Häusinger, Hung-Ni Chen, Lion Henrich, Hoa Ho, Friederike Reichel, Marcel Quint, Carla Mirabella, and Jinju Rhee. I want to thank Sabine Wilhelm-Kauf and all student assistants at the chair for their continuous support and for taking the time to chat with me to escape from research for a short while. I gratefully acknowledge funding from the German Research Foundation (DFG) through CRC TRR 190 and the GRK 1928.

Above all, I thank my friends and family for their support over the last years.

Contents

Contents	v
List of Figures	ix
List of Tables	xi
Preface	1
1 Comparison of Standard Auctions with an Agency Problem	5
1.1 Introduction	5
1.2 Model	10
1.2.1 Motivation and Background	10
1.2.2 Set-up	12
1.3 Analysis	14
1.3.1 No Punishment	14
1.3.2 Punishment	20
1.4 Discussion	24
2 Everyone Likes to be Liked – Experimental Evidence for Re-	
 ciprocal Preferences	27
2.1 Introduction	27
2.2 Experimental Design	33
2.3 Reciprocal Preferences and the Instability of Matching Markets	38
2.3.1 Instability of the Deferred Acceptance Mechanism	38

2.3.2	Reciprocal Preferences and Preference Adjustments . . .	39
2.4	Mechanisms Underlying Reciprocal Preferences	41
2.4.1	Theoretical Framework	41
2.4.2	Evidence for a Belief-Based Mechanism	45
2.4.3	Evidence for a Preference-Based Mechanism	48
2.4.4	Unconditional Cooperation	49
2.4.5	Mechanisms across Treatments	51
2.4.6	Gender Heterogeneity	52
2.4.7	Similarity, Homophily, and Reciprocal Preferences	53
2.5	Discussion	55
3	Reciprocal Preferences in Matching Markets	59
3.1	Introduction	59
3.2	Illustrative Example	65
3.3	Model	68
3.4	Two-Sided Matching Markets: Deferred Acceptance Mechanism	73
3.5	Stability in Two-Sided Matching Markets	77
3.6	School Choice Markets	80
3.7	Discussion and Conclusion	84
	Appendices	88
A	Appendix to Chapter 1	89
A.1	Proof of Proposition 1.2	89
A.2	Proof of Proposition 1.4	91
A.3	Additional Proofs	96
A.3.1	Manager's Bid greater than Bid in Standard Setting. . .	96
A.3.2	Monotonic Increasing Bidding Function	97
B	Appendix to Chapter 2	99
B.1	Behavioral Theory Model	99
B.1.1	Proof of Proposition 2.1	99

B.1.2	Proof of Proposition 2.2	101
B.2	Preregistered Analyses	103
B.3	Exploratory Analyses	106
B.4	Instructions	109
B.4.1	General Instructions (before Part I)	109
B.4.2	Questionnaire (Part I)	111
B.4.3	Instructions (Part II)	112
B.4.4	Additional Instructions (Part III)	124
C	Appendix to Chapter 3	127
C.1	Proof of Proposition 3.1	127
C.2	Proofs of Additional Statements	128
C.2.1	Statement 1	128
C.2.2	Statement 2	129
C.2.3	Statement 3	130
C.2.4	Statement 4	132
C.2.5	Statement 5	133
C.2.6	Statement 6	134
C.2.7	Statement 7	135

List of Figures

1.1	Timeline	13
1.2	No Punishment - SPSB	16
1.3	No Punishment - FPSB	18
1.4	Punishment - FPSB	23
2.1	Design Overview	34
2.2	Beliefs about Receiver's PGG Contributions: Proposers in <i>Info</i>	46
2.3	Average Conditional PGG Contributions: Proposer in <i>Info</i>	49
2.4	Unconditional PGG Contributions: Proposers in <i>Info</i>	50
3.1	Timeline	72
B.1	Contributing c^{max} by the Proposer	102
B.2	Unconditional PGG Contributions: Receiver	104
B.3	Beliefs of Receivers: PGG Contributions of Partner	108
B.4	Payoffs PGG: Implementation of Unconditional Decisions	109
B.5	Instruction - Proposer: Questionnaire	113
B.6	Instruction - Proposer: Preference Order	114
B.7	Instruction - Proposer: Adjustment of preferences	115
B.8	Instruction - Proposer: Contribution	117
B.9	Instruction - Proposer: Contribution table	118
B.10	Instruction - Receiver: Questionnaire	120
B.11	Instruction - Receiver: Preference order	121
B.12	Instruction - Receiver: Contribution	123

List of Tables

2.1	PGG Behavior of Proposers in <i>Info</i>	47
2.2	PGG Behavior of Proposers in <i>Info</i> and <i>No-Info</i>	52
2.3	Gender Heterogeneity of Proposers in <i>Info</i>	53
2.4	Homophily and Unconditional Contributions of Proposers in <i>Info</i>	55
3.1	Example 1	67
3.2	Example 2	76
3.3	Example 3	78
3.4	Example 4	81
B.1	Preference Adjustments across Treatments	103
B.2	Consistency of Preference Adjustments	104
B.3	Unconditional PGG Contributions of Receivers	105
B.4	Determinants of Proposers' Preference Adjustments	106
B.5	PGG Contributions and Questionnaire Responses	107
B.6	Beliefs of Receivers: PGG Contributions of Partner	108
B.7	Instruction - Loss attitudes	125
C.1	Example 5	128
C.2	Example 6	130
C.3	Example 7	131
C.4	Example 8	132
C.5	Example 9	133
C.6	Example 10	134

C.7 Example 11 135

Preface

“Market design is a kind of economic engineering, utilizing laboratory research, game theory, algorithms, simulations, and more. Its challenges inspire us to rethink longstanding fundamentals of economic theory.”

Paul Milgrom

Many markets, such as student assignment to schools or carbon trading, must be carefully designed to function efficiently. Market design is the economic discipline whose main objective is to create well-functioning markets by defining their rules.

Market design focuses on the behavior of market participants to ensure efficient market outcomes. The rules must be designed so that the behavior of individual market participants does not interfere with the desired market outcome. Participants may try to disrupt the market, take advantage of other participants, or exploit a loophole to maximize their own profit. Only if such behavior is prevented an efficient market outcome can be achieved, and more market participants may choose to enter the market.

Observing and analyzing existing markets helps to understand the incentives and actions of participants and thus the outcomes of markets. This knowledge can be used to further improve the design of these markets. The idea of this dissertation is to take relevant market observations, analyze them theoretically and experimentally, and thus contribute to a better understanding of these observations. This knowledge can then in turn be used to improve the design of markets.

This dissertation consists of one chapter about auction theory and two chapters about matching markets. Auction theory analyzes the optimal bidding strategy depending on the auction format and the information bidders have. The primary goal of matching market is to create stable matches, mostly between agents from two different sides of the market. Examples are the assignment of children to kindergartens or the allocation of teachers to schools.

Chapter 1: This chapter analyzes an auction setting where an owner of a company delegates bidding to a manager. The owner relies on the manager because only the manager is an expert in the market. The manager's incentives are only partially aligned with the owner. The manager has the incentive to maximize firm profits, but also gains utility from solely winning the auction because he has career concerns. The paper's primary research question is how this setup influences the optimal bidding strategy of a manager in the standard First Price (FPSB) and Second Price (SPSB) auctions. The analysis shows that there is an incentive to overbid in both the FPSB and SPSB auctions. If the owner can punish the manager in the case of a negative profit and does so, the expected price in the FPSB auction is higher than in the SPSB. These results may explain the observation that in many complex auction situations, auctioneers choose FPSB auctions.

Chapter 2: In this chapter, we analyze the observation that individuals often prefer to interact with those who want to interact with them. We investigate the effect of such "reciprocal preferences" on matching markets in a laboratory experiment. Matching markets can be unstable when individuals prefer to be matched with a partner who also likes them. We provide evidence that reciprocal preferences exist through a pre-registered and theory-guided laboratory experiment. Individuals, in fact, prefer to be matched with someone who rates them well. We show that this substantially decreases stability in matching markets and investigate underlying motives of reciprocal preferences.

Chapter 3: The second paper on reciprocal preferences studies their effect on stability and truth telling in a theoretical model. We formalize reciprocal preferences, apply them to matching markets, and analyze implications for mechanism design. We show that neither the common Deferred Acceptance mechanism nor any other mechanism is stable in standard two-sided markets. Observing the final allocation of the mechanism enables agents to learn about each other's preferences, which leads to instability. These results contribute to the understanding of non-standard preferences in matching markets and their implications for efficient information and mechanism design.

All chapters focus on topics that are, on the one hand, relevant in real-world settings, on the other hand, fairly new to be studied. Consequently, they help explain frequent observation in those markets that have been difficult to explain by standard economic reasoning. For example, it is widespread that managers at companies are responsible for bidding in an auction because of their expertise. At the same time, FPSB auctions are often implemented in those markets. Chapter 1 connects both observations by showing that the expected price of a FPSB auction is higher in such a setting than in a SPSB auction. Hence, an auctioneer prefers to implement a FPSB auction. In Chapters 2 and 3, the intuitive feeling that people like to be liked is combined with the observation that standard mechanisms are sometimes adapted to accommodate reciprocal preferences. For example, Avery and Levin (2010) show that colleges use early admission programs to admit students who strongly prefer to attend that college. Chapters 2 and 3 study the origin of reciprocal preferences and their effect on matching markets and their stability.

Chapter 1

Comparison of Standard Auctions with an Agency Problem

1.1 Introduction

In many real world auctions, bidding is delegated to agents. For example, in procurement auctions the owner (she) of a company typically delegates the bidding to a manager (he) who is better informed about the market and the good's or project's valuation. Similarly, a house buyer will often delegate the bidding for a house to a real estate agent who knows the housing market better. If the manager is not residual claimant on profits and if the manager gets an additional benefit from winning the auction (e.g., due to reputational or career concerns), preferences of the owner and the manager are misaligned. Successfully buying a house could signal to new potential clients that the real estate agent is of high quality. Winning a significant project might increase the reputation of a manager and improve the chance of a career within and outside of the company. The problem of a manager's empire-building is consistent with this idea (Jensen, 1986). Managers have incentives to expand their sphere of

influence beyond the optimal size to increase their resources and power. This paper studies the effect of this organizational owner/ manager setting and the misaligned incentives on the optimal bidding strategy.

The empirical auction literature highlights the widespread use of First Price Sealed Bid (FPSB) auctions in complex settings. In their chapter for the Handbook of Industrial Organization, Hortaçsu and Perrigne (2021) state that "projects and services are allocated through first-price sealed-bid procurement auctions or scoring auctions" (p.84), "most procurement auctions are conducted as sealed-bid first-price (lowest bid) auctions" (p.101), "oil and gas lease are usually sold through first-price sealed-bid auctions" (p.124). This paper provides one possible answer for the widespread use of FPSB auctions in such settings.

Throughout the analysis, I assume a standard independent private value auction setting with symmetric and risk neutral bidders. The focus is on the distinction between owner and manager and their different utility functions. The manager's incentive is modeled as career concerns that increase with the profit made and that yield a one-time fixed benefit if the auction is won.¹ In contrast, the owner's utility depends only on the profit. All bidders share the same organizational setting as an owner and a manager. The paper analyzes the setting in the four standard auction types: FPSB, SPSB, English, and Dutch auction. As in the standard setting, the FPSB and Dutch auction are strategically equivalent and the SPSB and the English auction are strategically equivalent. Hence, all results for the FPSB auction hold for the Dutch auction, and all results for the SPSB auction hold for the English auction.

In the first step, I assume that the owner cannot offer any incentive scheme to affect the managers's bidding behavior. The manager is motivated by reputational and career concerns only. His utility is composed of a constant fraction of the owner's monetary payoff plus an additional utility of winning the auc-

¹The approach and idea are related to the literature on the joy of winning (Cooper & Fang, 2008). This literature assumes that a bidder receives a one-time utility bonus if he wins the auction. The joy of winning is used to explain auction fever/ overbidding, e.g., in eBay auctions.

tion. The owner does not or cannot influence the bidding of her manager in any way. The situation is realistic in some settings and helps to understand the model's intuition and general properties. Due to the additional benefit of winning, a manager is incentivized to bid higher than in the standard symmetric equilibrium and compared to what the owner wants him to bid. The incentive to overbid is present in the FPSB and the SPSB auction. In this paper, overbidding is defined as a bid that is higher than the optimal bid in a standard setting where there are no additional incentives for the manager. Therefore, overbidding does not always imply that a bid is higher than the valuation. In both auction types, the manager bids according to a bidding function equal to the bidding function in a standard setting plus a fixed amount. The fixed amount is independent of the valuation and the auction type. The expected price in both auctions is, therefore, identical. However, the FPSB and SPSB auctions differ with regards to the risk of a negative profit. In a SPSB auction, the bid is always higher than the valuation. Hence, the profit can be negative for every valuation depending on the second highest bid. In a FPSB auction, the manager bids higher than the valuation only for small valuations.² He still overbids, but bids less than the valuation for high valuations. Therefore, the profit is only negative if a bidder with a small valuation wins.

In the second step, I assume that the owner can punish the manager if the manager's bid results in a negative profit for the owner. The manager must ensure that the auction's profit is never negative in order not to be penalized. If the punishment is reasonably harsh, the manager avoids bidding higher than the valuation in both auction types. In a SPSB auction, the manager bids exactly the valuation and does not overbid. A bid higher than the valuation always incurs the risk of a negative profit. In a FPSB auction, a manager avoids bidding higher than the valuation. A bidder with a value lower than a threshold bids exactly the valuation because he is primarily interested in maximizing the

²If the career concern mainly depends on winning the auction and not on the profit made, a manager bids higher than the valuation for every realized valuation.

probability of winning. For valuations higher than the threshold, he bids less than the valuation. However, the manager still bids higher than the optimal bid in a standard setting. Hence, with valuations higher than the threshold, he makes a positive profit if he wins, but the profit is lower than in the standard setting. As a result, while the manager will not overbid in a SPSB auction, he will overbid in a FPSB auction. Hence, the expected price and profit for the auctioneer is higher in a FPSB auction, and the auctioneer prefers a FPSB auction over a SPSB auction.

The main idea of why the FPSB auction generates a higher price than a SPSB auction is related to Che and Gale (1998). The authors show the effect of budget constrained bidders on the optimal bidding strategy. They find that the expected price is higher in a FPSB auction. Similar to the main idea of this paper, the constraint (in their case, a budget constraint) is less restrictive in a FPSB auction than in a SPSB auction. Since the optimal bid in a FPSB auction is lower than in a SPSB auction, the budget constraint in a SPSB auction is binding for more bidders than in a FPSB auction. This paper's owner/ manager setting contributes to the nascent literature on organizational structures in auctions. On a general level, this literature analyzes the impact of an organizational environment with misaligned incentives between an owner and a manager. Burkett (2015) focuses on the optimal budget constraint an owner should set up to prevent the manager from overbidding. Malenko and Tsoy (2019) consider a setting in which an advisor only counsels a buyer, but the buyer still bids in the auction. Both papers consider an owner who is well informed about the market but does not know the object's precise value. Conversely, this paper examines an owner who lacks information about general market properties and his valuation for the good. In general, market information is required to make a sophisticated bid.³ For example, Burkett (2015) assumes that the owner has a

³Bidding in a SPSB auction is an exception to this in the case of private values. If valuations are common or affiliated, the optimal bidding strategy in a SPSB auction also relies on market information like the number of bidders or the distributions the valuations are drawn from.

signal about the good's value that is not as informative as the manager's signal. While Burkett (2015) focuses on the owner's ability to influence the manager's bidding through a budget constraint, this paper focuses on the possibility of influencing the bidding by observing the profit and punishing the manager in case of a negative profit. Burkett (2015) finds no difference in revenue and efficiency between FPSB and SPSB auctions for independent valuations. In the case of affiliated valuations a solution can only be partially characterized. The FPSB auction is more efficient, while the effect on revenue is ambiguous. Malenko and Tsoy (2019) present a scenario in which a biased advisor (manager) consults an uninformed owner. The main difference in this paper is that the buyer controls the bidding decision. Hence, the theoretical analysis relates to a game of cheap talk. Malenko and Tsoy (2019) show that standard static auctions like FPSB and SPSB auctions lead to an equilibrium with communication taking an interval partition form. The distribution of possible valuations is separated into partitions. The advice within a partition always leads to the same bid by the owner. They show that all static auctions generate the same expected revenue. Furthermore, due to its dynamic nature, an English auction outperforms static auctions in terms of revenue and efficiency. Beyond that, Ausubel, Burkett, and Filiz-Ozbay (2017) test the model of Burkett (2015) in a laboratory experiment and find that the FPSB and SPSB auctions are equally efficient. Bichler and Paulsen (2018) analyze a similar Principal-Agent relationship in the context of spectrum auctions. They focus on multi-unit auctions.

Section 1.2 outlines a detailed motivation for the model's main assumptions by analyzing real world auction examples. It also sets up the model. Section 1.3.1 analyzes the model's implications if the owner cannot punish the manager for a negative profit. Section 1.3.2 relaxes this assumption and analyzes the effect of potential punishment. Section 1.4 discusses the implications of the model and concludes.

1.2 Model

Section 1.2.1 provides a short overview of empirical observations in auction settings. This background information serves as the basis for this study and motivates the assumptions in the model set-up in section 1.2.2.

1.2.1 Motivation and Background

This section highlights four main assumptions underlying this paper's model. First, managers may have an incentive to overbid, for example, due to career concerns or empire-building motives. Second, in many auction settings, only an expert can value the good or the project. Third, knowing the valuation of the good is insufficient, as a sound understanding of the market is essential for a sophisticated bid. Fourth, evaluating the profit of a good or project is challenging, and it is not always possible to infer from the profit whether the manager made a reasonable bid.

First, winning an auction might influence the manager's utility not only through the realized profit. Additionally, just winning the project might also be relevant for the manager. The empire-building problem is a well-known problem in the literature (Jensen, 1986).⁴ Managers may be incentivized to increase their department, team, and area of responsibility beyond the optimal size. These changes expand their power by increasing the resources they control and thus their salary. Additionally, winning an auction is a signal of success. It can positively influence the manager's career prospects within and outside of the firm, especially if the profit is unobservable or difficult to evaluate.

Second, only an expert can accurately assess the value of a good or project. For example, in timber or oil auctions, it is hard to estimate the value of a project without market experience. In timber auctions, tracts of forests are auctioned and bidders must conduct surveys to evaluate the timber volume of

⁴Jensen (1986) focuses on the relationship between a manager and an owner of a company, but the same argument can be made for managers on every level in a company.

each species in order to estimate their bid (Athey & Levin, 2001). Incomplete contracts and renegotiations are prevalent in construction projects. Renegotiations have a high potential for bidders to increase their profits later (Bajari, Houghton, & Tadelis, 2014). Knowledge about these potentials is crucial at the stage of the auction. In oil auctions, a bidder has to estimate the value after analyzing the geological condition. Precise knowledge about the regional conditions, extraction amounts of nearby tracts, and broad experience are indispensable (Hendricks, Porter, & Tan, 1993). Furthermore, reasonable arguments can be made supporting the notion that valuations are not independent but rather affiliated (or common values) in markets such as timber auctions, procurement of construction projects, oil and gas lease auctions, spectrum auctions, auctions for financial markets, etc. (Compiani, Haile, & Sant'Anna, 2020; Li & Zhang, 2010).

Third, while knowing one's valuation is an essential first step, a bidder can make a sophisticated bid only with sufficient market information. In most auction settings, the bidding function depends not only on the valuation but also on market information such as the number of bidders, risk aversion of bidders, bidder symmetry, and more.⁵ Only experts have sufficient market information to discern the potential valuation distributions and other relevant market information. In timber auctions, bidders can be distinguished between loggers and mills (Athey, Levin, & Seira, 2011). Loggers and mills are systematically different. In addition, sealed-bid auctions attract more loggers and generate more revenue. It is crucial to know whether auction competitors are loggers or mills. Knowledge of existing asymmetries is also relevant in construction projects. Bidders can systematically differ according to their distances to the contract location, job seniority and experience, and levels of risk aversion (Flambard & Perrigne, 2006). Therefore, knowledge about the market and other bidders is essential for successfully competing in the auction.

⁵The SPSB and English auction are an exception to this, but only in the case of private values.

Fourth, depending on the project, it is sometimes challenging for an owner to evaluate potential project profits due to long project durations and internal fixed costs allocations. In addition, uncertainties about fluctuating costs (e.g., wages, raw material), exchange rates, and other project characteristics (Jung et al., 2019; Kosmopoulou & Zhou, 2014) can further complicate a manager's bid quality assessment. Owners may try to avoid unfair punishment and only punish if they are sure a bid was unreasonable. This paper analyzes both a setting where the owner completely restrains from punishment (Section 1.3.1) and a setting where an owner punishes in the case of a negative profit (Section 1.3.2).

1.2.2 Set-up

An auctioneer wants to sell a single object. There are n bidders, $N = 1, \dots, n$, participating in the auction. Every bidder has a private value v_i for the object.⁶ Each private value v is drawn independently and identically from a uniform distribution $F(v)$ on $[0, 1]$ with density $f(v) = 1$. The value for the auctioneer is 0. Owners, managers and the auctioneer are all risk neutral. The winner of the auction gains his valuation v and pays the price p ; all other bidders gain nothing. The price p is either determined by the bidder's bid or by the bid of the second-highest bidder, depending on the auction format. The owner's profit is $\pi = v - p$.

The owner pays the manager a fixed wage \bar{w} , which is exogenously determined by the market. Without loss of generality, the manager's fixed wage is set to zero $\bar{w} = 0$ in the analysis. The manager's utility depends on \bar{w} and his career concerns. A successful auction helps the manager to make a career within the firm or to find a better job at a competitor. The manager receives a fixed utility a by just winning the auction. In addition, the benefit from career concerns increases with the profit $s\pi$. Factor s demonstrates the importance

⁶I omit the subscript i when possible and a distinction between owners (or managers) i and j is not necessary for understanding.




Figure 1.1: Timeline

of π for his career concerns. The simplest representation of career concerns is given by $h^s(a, \pi) = a + s(\pi) = a + s(v - p)$. The incentives a and s are the same for every manager and common knowledge to every player.

The manager knows v before the auction and is an expert in the market. He has all the relevant information about the market and knows that all managers' values are drawn from the distribution $F(v)$. Contrary to the manager, the owner does not have this information and therefore depends on the manager to bid in the auction. Figure 1.1 depicts the timeline. Each owner hires a manager to bid in the auction. The auctioneer decides on the auction design.

The relation between a and s plays a crucial role in the paper's analysis. The ratio a/s characterizes the manager's incentive of winning the auction relative to earning higher profits for the firm. The following two extreme cases illustrate this intuition. Assume that a manager is solely interested in winning an auction ($a > 0$ and $s = 0$). In this case, he is only motivated to win the auction but does not care about the profit. It is optimal for the manager to bid infinitely high, which results in a profit level of negative infinity for the owner. This outcome holds for both auction types (i.e. FPSB and SPSB auctions). In the other extreme case, the manager is only interested in maximizing the profit ($a = 0$ and $s > 0$). Hence, his preferences are perfectly aligned with the owner's preferences. He bids according to the standard symmetric bidding strategy without an additional incentive to win. His bid maximizes the owner's expected profit. These two extreme cases indicate that the manager tends to overbid even for small values of a and, therefore, small values of a/s .

The owner wants to punish a manager for overbidding. However, she also wants to ensure that she does not make an unfair punishment. An owner can only be sure that a negative profit is due to overbidding if the profit relies solely

on the manager's bid and not on some market uncertainties. If this is the case, a negative profit is a clear signal for overbidding. The owner then punishes the manager without the risk of unfair punishment. A setting without the possible punishment is analyzed in Section 1.3.1. Section 1.3.2 builds on the results of Section 1.3.1 and studies a setting with possible punishment.

The paper analyzes the delegation of bidding to a manager in the four standard auction types (i.e., FPSB, SPSB, English, and Dutch auction) in an independent private values setting. All results and proofs for the FPSB auction hold for the Dutch auction and all proofs for the SPSB auction hold for the English auction. Hence, as in the standard case, the FPSB and Dutch auction and the SPSB and English auction are strategically equivalent.

1.3 Analysis

This paper considers three different cases: 1) a setting without punishment, 2) a setting with punishment, and 3) a standard setting without an additional utility in case of winning. Section 1.3.1 analyzes a benchmark case where the owner does not punish a manager for a negative profit. Section 1.3.2 examines the manager's optimal bidding strategy if the owner can punish the manager. Both results are compared to the standard case. The setting is well known from standard auction literature, without agency problems.

1.3.1 No Punishment

The analysis shows that all managers overbid when they do not fear a punishment. All managers overbid by the same fixed amount, regardless of the type of auction and value v . Hence, the expected price in both auction formats is the same. An owner is rational and understands the situation he faces.

SPSB

In the standard case and a SPSB auction, it is a weakly dominant strategy for

a manager to bid exactly the value ($b(v) = v$). In this case, when the auction is won, the profit is always non-negative, and the price depends on the second-highest bidder. A similar logic applies to the setting of this paper. The manager has a weakly dominant strategy. Unlike in the standard setting, the optimal bid b depends not only on v , but also on a and s .

Proposition 1.1. *The optimal symmetric bidding strategy in a SPSB auction with n managers, all of whom have career concerns, is:*

$$b(v)_{NoPunishment}^{SPSB} = v + \frac{a}{s} \quad (1.1)$$

A manager is therefore overbidding by a/s compared to the standard setting.

Proof. Assume that a manager bids $b(v) = v + \frac{a}{s}$. By doing so, a manager's utility is always non-negative.⁷ He can deviate from his strategy in two ways. First, bidding less never increases his utility, but it can decrease it. In this case, either he wins the auction and gains the same profit, or he loses an auction he otherwise would have won with a non-negative utility. On the other hand, the manager might deviate by bidding higher than b . By bidding higher, he only wins in additional cases in which he receives a negative utility. However, he does not receive additional utility for cases he would have also won with bidding b . Therefore, bidding $b(v) = v + \frac{a}{s}$ is a weakly dominant strategy for every manager. \square

Figure 1.2 shows the difference between the optimal bid $b_{Standard}^{SPSB}$ in a standard auction setting and a manager's optimal bid $b_{NoPunishment}^{SPSB}$. For every value v , the manager bids a/s more than in the standard setting. The bidding function $b_{Standard}^{SPSB}$ also represents the bidding function that maximizes the profit and utility of the owner.

⁷If the second highest bidder j bids the identical amount $b(v_i) = v_i = b(v_j)$, his utility is $u_i = 0$ ($= a + s(v_i - c) = a + s(v_i - (v_i + \frac{a}{s}))$).

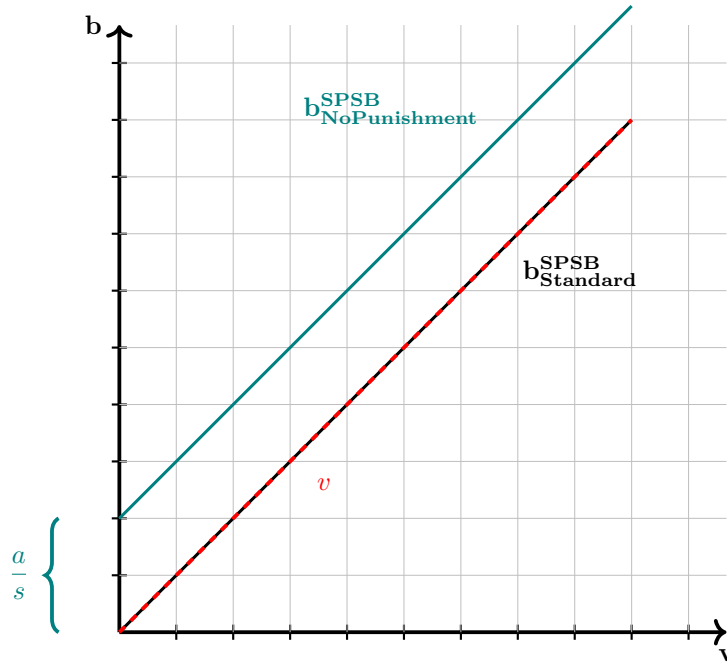


Figure 1.2: No Punishment - SPSB

FPSB

As in the SPSB auction, the manager has an incentive to overbid compared to the standard setting. The analysis is more complex than that of the SPSB auction. Due to the symmetric managers, the model's outcome is a Symmetric Bayesian Nash Equilibrium. In a standard setting, a bidder maximizes the expected profit $\pi(b_i) = (v_i - b_i) \cdot Pr(b_i > b_j)$. The bidder chooses the optimal strategy b_i depending on the bids of the other bidders. This equation illustrates each bidder's trade-offs. Increasing the bid increases the probability of winning while decreasing the profit conditional on winning and vice versa. Every manager with the organizational setting of an owner and manager faces the same problem. He maximizes his expected utility $\pi_i(b_i) = (s(v_i - b_i) + a) \cdot Pr(b_i > b_j)$ by taking a and s into consideration.

The optimal bidding function in the standard FPSB auction is $b_{Standard}^{FPSB}$. Compared to this standard setting, the manager's utility has an additional component a from winning the auction. Bidding $b_{Standard}^{FPSB}$ is no longer his best strategy. Because of the added utility a , he wants to increase his probability of

winning and accepts that the (monetary) profit is lower in the case of winning. He bids higher than in the standard case to increase his probability of winning. Every manager follows the same logic, thus resulting in higher bids.

Proposition 1.2. *The optimal symmetric bidding strategy in a FPSB auction with n managers, all of whom have career concerns, is:*

$$b(v)_{NoPunishment}^{FPSB} = \frac{n-1}{n}v + \frac{a}{s} \quad (1.2)$$

*A manager is therefore overbidding by a/s compared to the standard setting.*⁸

Proof in the appendix.

Proposition 1.2 and Figure 1.3 demonstrate overbidding by a manager by a/s . The bidding function is shifted upwards by a/s . For low values, the bid is higher than the valuation ($b > v$). Therefore the profit is negative in the case of winning. To understand the intuition for the shift by a/s , it helps to analyze a manager with the lowest possible value $v = 0$.

A bidder in the standard case with a valuation of $v = 0$ bids $b_{Standard}^{FPSB}(0) = 0$ and has an expected utility of 0. Similarly, a manager with $v = 0$ bids $b_{NoPunishment}^{FPSB}(0) = \frac{a}{s}$. If he would win, his utility is 0 ($= a + s(v - c) = a + s(0 - \frac{a}{s})$).⁹ Given a strictly monotonic increasing bidding function, the expected profit for a bidder with the lowest possible value $v = 0$ must be 0 due to the 0 probability of winning. If the optimal bid $\hat{b}_{NoPunishment}^{FPSB}(v = 0)$ would be smaller than a/s , the bidder with a value $v = 0$ could deviate from the strategy $\hat{b}_{NoPunishment}^{FPSB}(v = 0)$ and bid more to increase his utility. If he bids an amount between $\hat{b}_{NoPunishment}^{FPSB}(v = 0)$ and a/s , he gains a positive expected utility because the probability of winning is positive and his utility in case of

⁸The optimal bid in a standard setting with a uniform distribution is $b(v) = \frac{n-1}{n}v$.

⁹A common result in auction theory is that a bidder with the lowest possible valuation has an expected utility of 0 (Myerson, 1981).

winning is positive. Hence, a bidding function that is strictly increasing cannot have an optimal bid for a bidder with the lowest possible value that is smaller than a/s .

Next, I argue that given every manager bids at minimum a/s , the maximization problem is similar to a setting without the additional utility a in case of winning. Assume that the optimal bidding function consists of two parts; a fixed amount a/s that everybody bids and the function $\beta(v)$ that depends on the value v ($b_{NoPunishment}^{FPSB}(v) = \beta(v) + \frac{a}{s}$). A manager wants to maximize his utility $u_i = (s(v_i - \beta(v_i)) - \frac{a}{s} + a) \cdot Pr(b_i > b_j) = ((s(v_i - \beta(v_i)) - a + a) \cdot Pr(b_i > b_j) = (s(v - \beta(v_i)) \cdot Pr(b_i > b_j)$. Hence, the problem is similar to a problem where there is no additional incentive of winning, but the agent only receives a share s of the profit. Hence, $\beta(v)$ should be equal to $b_{Standard}^{FPSB}$ because $b_{Standard}^{FPSB}$ maximizes the utility given everybody else bids according to $b_{Standard}^{FPSB}$. This logic is independent of the given distribution function.




Figure 1.3: No Punishment - FPSB

In the next step, I compare the expected price in a FPSB and SPSB auction.

Lemma 1.1. *The auction's expected price when managers are incentivized by career concerns and face no punishment by the owner is the same for the FPSB and SPSB auctions.*

Proof. In both auctions, all bidders increase their bid by a/s compared to the equilibrium in a standard setting. In both auction types, the bidding function is equal to a standard auction setting with values independently and identically drawn from a distribution with support $\left[\frac{a}{s}, 1 + \frac{a}{s}\right]$. We know that the SPSB and FPSB auctions generate the same expected profit in this context. Furthermore, the auctioneer is indifferent between these auction types.¹⁰ \square

The auctioneer's profit in the setting with an agency problem on the bidder side is higher than in a standard setting. The owners suffer from the managers' overbids as they risk a negative profit. In contrast, the level of overbidding in this setting does not increase the manager's utility when compared to a standard setting. Due to the bidder's symmetry and the sales managers overbidding, no manager increases his probability of winning compared to the standard setting. While a manager receives an additional utility of a in case of winning, he overbids by a/s . Thus, the profit reduces by a/s . The negative effect on the utility of the manager due to the lower profit is a ($= (a/s) \cdot s$). Hence, both effects cancel out. The level of a does not influence a manager's utility.

Depending on the level of a/s , the consequences might be extreme. Overbidding would be a common phenomenon. Owners will restrain managers from participating in an auction if their expected profit is negative. The following section introduces a scenario in which the owner can punish the manager for a negative profit. The manager anticipates this and refrains from bidding higher than v .

¹⁰The revenue equivalence theorem (Myerson, 1981) shows that given standard assumptions the expected revenue of an auction is independent of the auction design. This paper shows that the FPSB and SPSB auctions have the same expected revenue given this setting.

1.3.2 Punishment

This section analyzes a scenario in which the owner can punish the manager for a negative profit. Due to the possibility of punishment, the manager restrains from bidding higher than the value. This affects the bidding function in the FPSB and SPSB auctions differently. The manager does not overbid in the SPSB auction but overbids in the FPSB auction. Hence, the expected price is higher in the FPSB auction. In many settings, it is reasonable that (after some time) the owner deduces the profit π of an auction/ project and uses it as a clear signal to determine if the manager overbid. Therefore, one way to influence bidding is to punish a manager if the owner identifies a negative profit.

The manager's utility u_m changes if the owner punishes him. In case of a negative profit, the manager suffers a negative utility z due to the owner's punishment.

$$u_m = \begin{cases} s(v - p) + a & \text{in case of winning the auction and if } \pi \geq 0 \\ s(v - p) + a - z & \text{in case of winning the auction and if } \pi < 0 \\ 0 & \text{else} \end{cases}$$

Any penalty of $z > a$ is sufficient for a manager to bid no higher than v in both auction types. Every punishment $z > a$ has the same effect on the manager's bidding. A punishment $z < a$ may not be strong enough to prevent managers from bidding higher than the valuation.

Suppose a manager participates in a FPSB auction and is not incentivized by profits ($s = 0$). He only cares about his career concerns a and would like to bid unlimitedly high. If he wins, his utility is negative $u_m = a - z$. Therefore he refrains from bidding higher than v because not bidding or bidding below v never generates negative utility. When he bids higher than v and $s > 0$, this negatively affects his utility because he suffers from negative profit. Therefore, $z > a$ is sufficient for the manager never to bid $b(v) > v$.

The results for the SPSB auction are similar; however, the underlying argument differs. Bidding $b_i > v_i$ instead of $b_i = v_i$ only affects the utility of a

manager i if the second-highest bid b_2 is between his bid b_i and the valuation v_i ($b_i > b_2 > v_i$). He then also wins in those cases. The profit is always negative because $b_2 > v_i$. Again, assume $s = 0$. In the case of $b_i > b_2 > v_i$, the manager receives additional utility a but loses utility z due to the punishment. If $z > a$, the manager never bids higher than v_i . If $s > 0$, there is a second effect that negatively impacts the manager's utility upon winning due to the negative profit.

SPSB

The analysis of a manager in a SPSB auction is straightforward. The manager wants to bid $b(v) = v + \frac{a}{s}$ without the punishment. However, the punishment prevents him from bidding $b(v) > v$. The optimal bid for a manager is $b(v) = v$.

Proposition 1.3. *The weakly dominant strategy in a SPSB auction with n managers, all of whom have career concerns and receive punishment $z > a$ in the case of a negative profit, is:*

$$b(v)_{\text{Punishment}}^{\text{SPSB}} = v \quad (1.3)$$

Proof. A manager can again deviate from strategy $b(v) = v$ in two ways. As shown before, if a manager bids higher than v ($b(v) > v$) the profit π is negative in the cases he additionally wins. Hence, he receives a negative utility due to the punishment. If he bids less than v ($b(v) < v$) he could miss out on cases in which he would have gained positive utility. Hence, it is a weakly dominant strategy to bid $b(v) = v$. \square

If he bids $b(v) = v$, the owner never makes a negative profit, and the manager never gets punished. Moreover, the manager bids as the owner wants him to do. This maximizes the owner's expected profit.

FPSB

Similar to the setting without punishment, the analysis of the FPSB auction is

non-trivial. Compared to the SPSB auction with a weakly dominant strategy for each player, the equilibrium in a FPSB auction is a Symmetric Bayesian Nash Equilibrium. Thus, each bidder's strategy affects the optimal bidding strategy for every other bidder.

To understand the bidding strategy, consider first the optimal bidding strategy without punishment. In this setting, the bidding function is shifted upwards compared to the bidding function in the standard setting. Given moderate values of a , this shift results in a negative profit for low values and a positive profit for high values. A manager with small values is restricted and cannot bid as high as in the setting without punishment. Bidders with high values could follow the optimal bidding strategy of the setting without punishment. They will make a positive profit and will not be punished. However, as bidders with low valuations are forced to bid less aggressively, bidders with high valuations adjust their bidding strategies and also bid less aggressively. The bidding strategy consists of two cases. For low values, managers bid their valuation $b(v) = v$ up to a threshold \bar{v} . For values higher than the threshold, managers bid less than the valuation but more than in the standard setting.

Proposition 1.4. *The optimal symmetric bidding strategy in a FPSB auction with n managers, all of whom have career concerns and receive punishment $z > a$ in case of a negative profit, consists of two parts. The threshold value is given by $\bar{v} = (n - 1)\frac{a}{s}$.*

$$b(v)_{Punishment}^{FPSB} = \begin{cases} v, & v \leq \bar{v} \\ \frac{n-1}{n}v + \frac{a}{s} - \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{nv^{n-1}}, & v > \bar{v} \end{cases}$$

Proof in the appendix.

Lemma 1.2. *The bidding function $b(v)_{Punishment}^{FPSB}$ in a FPSB auction and managers incentivized by career concerns is monotonic increasing and always higher than the bidding function $b(v)_{Standard}^{FPSB}$ in the standard setting.*

Proof in the appendix.




Figure 1.4: Punishment - FPSB

Up to threshold \bar{v} , the additional career concerns a are high enough that a manager maximizes his utility through increasing the probability of winning and accepting a zero profit. Therefore, he bids as high as possible without being punished. He overbids compared to the equilibrium in the standard setting (see Figure 1.4). Managers with a value greater than \bar{v} bid lower than the valuation v and make a positive profit in the case of winning. This results in a positive profit for the owner. Still, the manager is bidding higher than in the standard setting ($v \geq b^{FPSB}_{Punishment} > b^{FPSB}_{Standard}$). For all positive values of v, a, s and n , the manager bids higher than the owner desires (see Appendix A.3.1).

Mathematically, the threshold can be greater than 1 ($\bar{v} > 1$). If this is the case, all managers bid equal to the value $b(v) = v$ independent of their valuation. The threshold value \bar{v} increases with high competition (n), high career concerns

(a), and low interest in profit (s).

Proposition 1.5. *The auction's expected price is higher in a FPSB than in SPSB auction when managers are incentivized by career concerns and there is a possibility of punishment by the owner.*

Proof. The possibility of punishment has different effects on the FPSB and SPSB auctions. Both auctions prevent the manager from bidding higher than the valuation. Hence, in a SPSB auction, a manager does not overbid. In the FPSB auction, a manager overbids because his bid is higher than the optimal bid in the standard setting. The manager can overbid without getting punished by the owner. The profit would be the same in both auction types if the optimal bidding function were similar to the bidding function in the respective standard setting. The bidding function in the FPSB auction is always higher than in the standard setting. Hence the expected price is higher in a FPSB than in a SPSB auction. \square

The FPSB auction generates higher revenue for the auctioneer and, therefore, is preferred by him. On the contrary, both owners and managers prefer a SPSB auction over a FPSB auction if $s > 0$ and $a > 0$. The manager's preference seems rather surprising. The manager's utility depends on the profit and the utility in case of winning a . Given v , the probability of winning is the same in both auction types due to symmetry between bidders. However, the expected profit is higher in a SPSB auction.

1.4 Discussion

There are two primary motivations for this paper. First, this study analyzes auction outcomes when the manager's interests are not fully aligned with the owner's preferences. Second, it presents one reason for the widespread use of FPSB auctions if the seller can choose the auction format. Both motivations are found in complex real world auction settings (e.g., timber, oil, construction, and

many more). The analysis shows that without any control by the owner in the form of punishment, overbidding is present in both FPSB and SPSB auctions. However, if the owner can punish a manager for negative profits, the effect differs between auction types. In both auction types, the threat of punishment prevents a manager from bidding higher than the valuation. Hence, a manager is bidding equal to the valuation in a SPSB auction. In a FPSB auction, however, the manager bids higher than in a standard setting but still less or equal to the valuation. Thus, the expected price is higher in the FPSB auction, and the auctioneer prefers a FPSB auction.

It is interesting to explore the case where career concerns are endogenous. The result with punishment in Section 1.3.2 highlights that managers prefer a SPSB auction over a FPSB auction given career concerns a . The auctioneer decides the auction format, and it is hard for a manager to influence that decision. In contrast, managers might be able to influence the career concerns structure. For example, winning an auction is only taken as a signal of success due to the lack of transparency in the observation and interpretation of the profit. Managers could reduce this informational advantage. This would result in a higher focus on the gained profit (higher s) and less on whether an auction was won or not (lower a). Hence, what level of a does a manager prefer given a FPSB auction? Showing that a high level of a is beneficial for a manager suggests that career concerns persist over time. The results in Section 1.3.1 show that the level of a has no effect on the utility of a manager in a FPSB auction. The positive effect of a higher a and the negative effect of a reduced profit due to overbidding cancel each other out for a manager. In the setting with punishment, the manager's overbidding is restricted. An increase of a has a weaker effect on the level of overbidding than in a setting without punishment. Given a value v , the optimal bid for a manager with punishment is always lower than in the setting without punishment ($b_{NoPunishment}^{FPSB} > b_{Punishment}^{FPSB}$).¹¹ Comparing the setting with and without punishment, an increase of a still has

¹¹Except for a bidder with value $v = 0$

the same positive effect on the manager's utility, but the negative effect of reduced profits is smaller in the setting with punishment. The manager cannot overbid by the same amount as in the setting without punishment. Hence, a higher level of a is better for a manager given the possibility of punishment.

One natural follow-up is to analyze the optimal contract from the owner's perspective. It seems reasonable that the first-best contract for an owner to offer is a contract that includes a monetary penalty in case of winning the auction that is equal to the one-time utility of winning a . This aligns both the manager and owner's preferences. However, this option seems unreasonable given a standard principal-agent model in which the principal wants to motivate the agent with a high salary to provide high effort. Most sales contracts pay a bonus in case of success (Joseph & Kalwani, 1998). Managers must work hard in order to be successful at an auction and the owner wants to incentivize the manager to exert this extra effort. Considering both effects, what is the best contract an owner can offer a manager, and what are other options for influencing her manager's bids?

Chapter 2

Everyone likes to be liked – Experimental Evidence for Reciprocal Preferences¹

2.1 Introduction

We often prefer to interact with individuals who also want to interact with us. For example, applicants may reconsider a job offer after learning they were not the first-choice candidate.² They may realize that an employer who does not favor them will be less invested in their relationship, or they may be less willing to invest in such a relationship themselves. We say that individuals who prefer to be matched with a partner who wants to be matched to them have *reciprocal preferences*.

Reciprocal preferences are especially relevant in matching markets because they impair efficiency. In many markets, participants (such as schools³) want

¹This chapter is based on joint work with Timm Opitz.

²See <https://www.forbes.com/sites/lizryan/2018/01/20/im-the-second-choice-candidate-should-i-still-take-the-job>, accessed 07/18/2022.

³Concerns were raised when it came to changes in the school admission system that had left principals uninformed about students' rankings of the schools (see <https://www.nytimes.com/2004/11/19/education/council-members-see-flaws-in-schooladmissions-plan.html>, accessed 07/18/2022).

to know others' preferences before submitting their ranking to the mechanism. Because preferences are typically not disclosed, existing mechanisms may be modified to incorporate them nonetheless. These changes can lead to inefficient market outcomes. For example, it became common for German universities to only accept medical students who had ranked the respective university favorably until the Federal Constitutional Board prohibited such practices.⁴ Early admission represents a more subtle form for universities to admit highly interested students. Avery and Levin (2010) show that “early admissions programs give students an opportunity to signal this enthusiasm.” Opitz and Schwaiger (2022b) show theoretically that reciprocal preferences impair stability in centralized matching mechanisms. For this, it is enough that agents observe the final matching and update their beliefs about others' preferences. This may lead agents to break up their match, contradicting the main objective of matching markets to establish stable relationships. In decentralized markets, reciprocal preferences may make it impossible for the decision-maker to fill a vacancy (Antler, 2019). Although anecdotal evidence and theory suggest that reciprocal preferences affect efficiency in matching markets, we cannot isolate them as the cause, because true preferences are not identifiable in observational data.

In this study, we identify reciprocal preferences and their impact on matching markets through a laboratory experiment. The experimental setting allows us to manipulate participants' information sets. We test whether agents' preferences are sensitive to information about others' preferences and analyze preference changes associated with this information. We quantify their impact on the stability of matching markets and differentiate underlying mechanisms. We hypothesize that participants prefer a partner who ranks them favorably (Aronson & Worchel, 1966; Montoya & Horton, 2012). Therefore, participants change their preference order after learning how others ranked them, which leads to instability in the matching market.

⁴For the ruling in Germany see *BVerfG, 1 BvL 3/14* (2017). Many institutions have similar procedures, such as Trinity College in Toronto, which only accepts students who rank the college first.

In the experiment, participants form two-person teams for a Public Goods Game (PGG) through a centralized matching mechanism. During the team-formation stage, participants interact in groups of eight, split equally into two market sides. Based on personality questionnaire responses, participants indicate with whom from the other market side they would like to play the PGG. They submit a rank-ordered list of potential partners from the other market side to a centralized Deferred Acceptance (DA) mechanism. The DA mechanism theoretically achieves stable allocations in two-sided matching markets (Gale & Shapley, 1962). In our treatment (*Info*), one side of the market receives information about with whom they are tentatively matched and how their potential partners rank them. In our baseline (*No-Info*), this market side never learns how their potential partners rank them and only sees with whom they are tentatively matched. In both treatments, they can subsequently change their preference list, resubmit it to the mechanism, and may get a new partner as a result. Afterward, the matched partners play a standard PGG, which captures the essential trade-off between collectively beneficial but individually costly contributions to a public good. This design allows us to understand the effects of reciprocal preferences on the stability of matching markets and investigates the underlying mechanisms.

We develop a stylized behavioral model to study two mechanisms for the emergence of reciprocal preferences in cooperative settings. The first channel is belief-based. It assumes that agents expect partners who like to be matched with them to be more cooperative (i.e., they expect their partner to contribute more in the PGG). The belief that favorable preferences signal a higher match-specific payoff provides a profit-maximizing rationale for reciprocal preferences. The second channel is preference-based. This channel posits that agents are more altruistic when matched with those who like them. As a consequence, they prefer to be more cooperative themselves (i.e., they contribute more in the PGG). Both channels imply that being matched with a partner who ranks the agent favorably spurs a higher utility, thereby providing a foundation for

reciprocal preferences. This experiment allows us to test both channels.

Our outcome variables are derived from the team-formation stage and the PGG. The first set of outcome variables investigates the effect of reciprocal preferences on stability. Achieving stable outcomes is key to matching mechanisms and implies Pareto efficiency (Gale & Shapley, 1962); Opitz and Schwaiger (2022b) show that reciprocal preferences can lead to instability when agents update their beliefs about others' preferences after the allocation of the mechanism.⁵ We analyze whether participants change their preference order once they learn how they are ranked by their potential partners, whether these preference changes are indicative of reciprocal preferences, and how this affects stability. The second set of outcome variables is based on subsequent behavior in the PGG and sheds light on belief-based and preference-based mechanisms underlying reciprocal preferences. We test whether reciprocal preferences are belief-based by analyzing incentivized beliefs about the partner's contributions to the PGG. To test the preference-based mechanism, we focus on conditional contribution decisions. In these decisions, we isolate altruism from the beliefs about a partner's contribution.

The main results can be summarized as follows: First, agents adjust their preferences significantly more often when they observe their potential partners' preferences (*Info*) than when they do not (27.67 vs. 9.67 percent). We find that preference adjustments in *Info* are consistent with reciprocal preferences. Participants rank those who rank them favorably higher than those who do not - they *like to be liked*. Second, these preference adjustments translate into significantly more unstable matchings in *Info* than in *No-Info* (40.00 vs. 10.67 percent). This outcome provides strong evidence that reciprocal preferences can inhibit the desired functioning of matching mechanisms. Third, our results indicate that both belief-based and preference-based motivations underlie

⁵Updating can either happen through directly learning others' preferences (as in the experimental design), or more subtly through observing the final matching and being able to make inferences about the underlying preferences that led to the matching. For a detailed analysis, see Opitz and Schwaiger (2022b).

reciprocal preferences. We show that participants hold (accurate) beliefs that someone who likes to be matched with them will be more cooperative. In this sense, revealed preferences signal the future value of the match, providing a profit-maximizing rationale for working with someone who likes you. In addition, we find evidence that participants act more altruistically towards those who indicated a preference towards them, providing support for a preference-based foundation. Lastly, the possibility of incorporating the potential partners' preferences into decision-making results in higher average cooperation and profits.

Our findings contribute to a better understanding of matching markets, team formation, and team behavior. First, we contribute to the growing experimental literature on matching markets (see Hakimov & Kübler, 2021, for a review). This literature attempts to uncover factors that limit the efficient functioning of matching markets because they affect agents' strategies. We are the first to study reciprocal preferences experimentally, and investigate whether outcomes of the DA mechanism remain stable. In this way, we test the empirical stability of the DA mechanism when others' preferences are revealed. Closest to this is previous work on the impacts of information about other participants' preference profiles in one-sided (Pais & Pintér, 2008) and two-sided centralized markets (Pais, Pintér, & Veszteg, 2011) on truth-telling. While these papers center on the extent to which agents use additional information to misrepresent their preferences strategically across mechanisms, we are interested in the causal effect of knowing one's rank in the preference order of potential partners on stability.

Second, we contribute to social preferences, team formation, and cooperation literature. Individuals often prefer to interact and team up with agents who are similar to them, which is known as *homophily* (McPherson, Smith-Lovin, & Cook, 2001). Homophily can be observed in experimental settings (R. Chen & Gong, 2018; Currarini & Mengel, 2016), as well as economic settings (i.e., the choice of co-workers and entrepreneurial teams (Boss et al., 2021; Hedegaard

& Tyran, 2018)). Self-selected teams display higher satisfaction, collaborative spirit, and effort (Boss et al., 2021; R. Chen & Gong, 2018), while results on performance are mixed.⁶ We contribute to the organizational literature on efficient team formation by highlighting the role of reciprocal preferences. We show that for an individual not only the similarity with their potential team partners matters, but also their partners' preferences.

Moreover, individuals are more likely to cooperate with those they perceive as similar and those who are generous towards them. An individual's prosocial behavior depends on group characteristics and behavior. People are more cooperative if they perceive others to belong to the same group (Akerlof & Kranton, 2000), where social identity may either be fostered through previous interaction (Eckel & Grossman, 2005), or by a shared preference (e.g., Y. Chen & Li, 2009, with a minimal group paradigm). Consequently, social proximity can overcome market imperfections (Chandrasekhar, Kinnan, & Larreguy, 2018; Jain, 2020), leading to higher levels of altruism (Goeree et al., 2010; Leider et al., 2009). Conversely, mutual dislike often hinders team behavior (Gerhards & Kosfeld, 2020). Our paper investigates one mechanism underlying these findings and causally identifies the effect of others' preferences on behavior. People treat those who have been generous towards them more favorably (Akerlof, 1982). We extend the recent literature on reciprocity towards non-monetary gifts (Bradler et al., 2016; Kube, Maréchal, & Puppe, 2012). We show that interpersonal preferences are another currency of reciprocity, most closely related to the idea in R. Dur (2009) that "employees care more for their manager when [...] their manager cares for them".

The paper is structured as follows: Section 2.2 presents our experimental design, Section 2.3 outlines our hypotheses and results on reciprocal preferences and the matching stage. Section 2.4 illustrates the underlying mechanisms through a stylized model, and investigates these mechanisms empirically.

⁶See Horwitz and Horwitz (2007) for a broader discussion on homophily and (workplace) performance.

Finally, we discuss and conclude in Section 2.5.

2.2 Experimental Design

Research Questions Through our experimental design, we examine three main research questions. First, do participants display reciprocal preferences? Second, do reciprocal preferences lead to instability in matching markets? Third, what are the mechanisms underlying the change in stated preferences? To address these questions causally, we exogenously manipulate information structures between treatments. This provides us with the necessary variation that observational data cannot give us to identify reciprocal preferences, their underlying mechanisms, and their implications for matching markets.

Overview The pre-registered experiment consists of three main parts.⁷ Part I collects self-reported personality data. In Part II, participants form two-player teams through a centralized matching mechanism and play a PGG within the formed dyads. Participants indicate with whom they would like to team up based on their potential partner’s personality profiles from Part I. In Part II, we compare behavior under two information structures in a between-subject design. In the treatment condition (*Info*), participants on one side of the market learn how their potential partners ranked them before submitting their final preference ranking. On the other hand, in *No-Info*, participants never know how their potential partners ranked them. Part III elicits beliefs about the PGG contribution of their team partner and collects control variables. The design is visualized in Figure 2.1.

Part I All participants fill out a personality questionnaire with 15 items on a four-point Likert scale. It contains five statements each on personality traits, preferred leisure activities, and societal opinions (see Appendix B.4.2 for the

⁷The preregistration of our design, as well as a detailed pre-analysis plan can be found at AEARCTR-0007551.

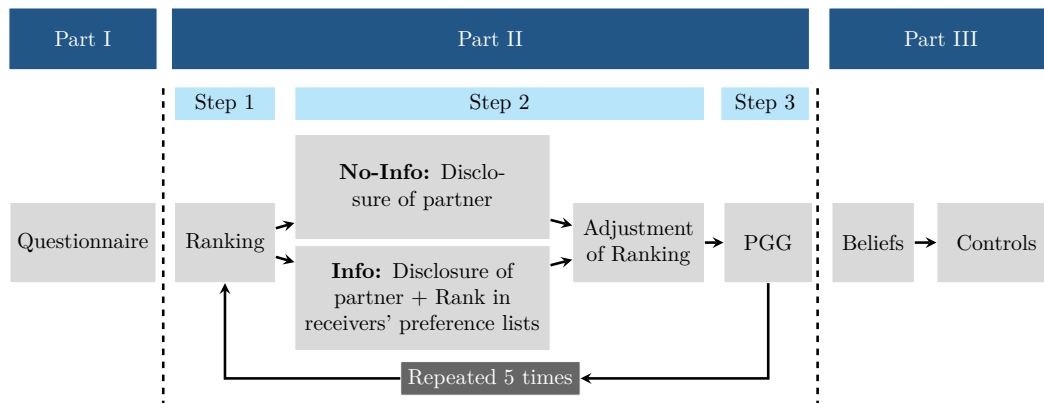


Figure 2.1: Design Overview

complete questionnaire).

Part II Participants are randomly assigned to one of two roles. Half take the role of proposers, half the role of receivers.⁸ Four proposers interact with four receivers in each *matching market*. The matching mechanism forms four teams, consisting of one proposer and one receiver. This procedure is the same in every *matching market*. Part II consists of three steps.

Step 1 Proposers and receivers submit a rank-ordered list of their potential partners. Proposers rank the four receivers in their *matching market* according to the desirability to be matched with them and vice versa. Teams are tentatively formed through the centralized DA mechanism (Gale & Shapley, 1962), which matches one proposer with one receiver for the upcoming PGG.⁹ Participants submit their preferences based on questionnaire responses from Part I. Each proposer in the *matching market* sees the same five randomly chosen answers from each receiver. The receivers see the answers to five different questions randomly selected among the remaining ones.¹⁰ After the participants submit

⁸In the beginning of Part II, participants are informed about their role, and receive detailed instructions on the procedures of the team-formation process and the PGG (see Appendix B.4).

⁹This means that we study a setting of two-sided matching in a one-to-one matching market, often referred to as a *marriage market* following Gale and Shapley (1962).

¹⁰The intuition for sharing distinct questions is to minimize the initial correlations between

their rank-ordered lists, the DA mechanism implements the tentative allocation.

Step 2 Proposers can submit a revised preference list to the DA mechanism. We vary the information between our two treatments *Info* and *No-Info*. In *No-Info*, proposers see with whom they have been matched in the first step. In *Info*, a proposer receives additional information on how all receivers ranked him. After examining this information, proposers decide on whether to revise their preference list and re-submit it to the DA mechanism. Receivers do not play an active role in this step as their preferences remain fixed. Furthermore, they do not receive any information about the proposer's actions during this step of the experiment. Proposers know that receivers never learn about proposers' preferences (and changes thereof).

Step 3 The formed dyads play a two-player PGG in the final step of Part II. Both partners receive an initial endowment of 10 Taler (experimental currency) and can either allocate it to a private account or contribute to a public account. The contributed amount of each partner $c_i \in \{0, 1, \dots, 10\}$ is referred to as the unconditional contribution. The sum of both players' contributions to the public good is multiplied by 1.5, and divided equally between the two. This leads to the following payoff function for a participant i : $\pi_i = 10 - c_i + 0.75 * (c_i + c_j)$. The marginal per capita return of 0.75 implies that free-riding ($c_i = 0$) is the dominant strategy from an individual perspective. However, since the sum of marginal returns is greater than 1, contributing the entire endowment of 10 Taler maximizes the team surplus. In addition to the unconditional contribution, proposers also fill in a table indicating their contribution for every possible contribution of their matched partner, referred to as their conditional contributions (Fischbacher, Gächter, & Fehr, 2001). Receivers only state their

preferences. If similarity is a relevant determinant for the choice of a partner (*homophily*), different questions provide different information about similarity, which reduces the correlation of preferences. In the extreme of perfect correlation, everyone is already matched with the partner they prefer most and that prefers them most, such that *reciprocal preferences* do not affect the outcome.

unconditional contribution.¹¹ The final payoff for the receiver depends on the stated unconditional contributions of both players. The final payoff for the proposer depends on the receiver’s unconditional contribution and on the proposer’s conditional or unconditional contribution.

Part III We complement the contributions to the PGG with incentivized point beliefs about partner’s unconditional contribution, for both proposers and receivers (Gächter, Kölle, & Quercia, 2017). We do not announce the belief elicitation before, to rule out that expectations about the ability to judge the behavior of another player influence preference submission.

We also elicit proxies for cognitive ability, loss attitudes, and socio-demographic controls. Proposers with higher cognitive abilities may be more likely to perceive receivers’ preferences as signals for their contribution and adjust their preferences strategically. To control for this, we use Raven’s Matrices.¹² Participants are given 5 minutes to complete increasingly difficult Raven’s Matrices, scored on the number of correct answers minus the number of incorrect answers. High degrees of loss aversion may make participants less likely to adjust their preferences if they feel attached to their current partner. Although unlikely given the information sets of participants in our experiment, (expectation-based) loss aversion may influence initial reporting strategies (Meisner & von Wangenheim, 2021). Hence, we elicit an incentivized measure of loss aversion in risky choices (Gächter, Johnson, & Herrmann, 2022). Before concluding the experiment, participants complete a short socio-demographic questionnaire.

¹¹This circumvents the problem with conditional contributions that the standard (unique) Nash-Equilibrium of not contributing anything requires common knowledge of rationality (Fischbacher, Gächter, & Fehr, 2001, Footnote 6). In light of a substantial fraction of conditional cooperators in previous PGG experiments, we do not want to assume this and let receivers only make an unconditional contribution decision (which is known to the proposers).

¹²The Raven’s Matrices test is a leading non-verbal measure of analytic intelligence, test scores are associated with the degree of sophistication in the beauty contest (Gill & Prowse, 2016), in manipulable matching mechanisms (Basteck & Mantovani, 2018), as well as with more accurate beliefs (Burks et al., 2009).

Repetitions We repeat Part II five times. During each repetition, participants play within a new *matching market* of randomly selected participants. Roles as proposer or receiver remain constant across rounds. To minimize the influence of earlier rounds on later rounds, participants do not receive feedback between rounds. Furthermore, by displaying only a subset of questionnaire responses in each round and randomly assigning participants to *matching markets*, we minimize the possibility that participants may identify others across rounds.

Payoffs and Incentive Compatibility One round of the PGG is randomly chosen to be payoff relevant. Participants earn money through their final payoff from the PGG (determined by their own and their partner’s contribution choice) in one of the five rounds. For proposers, we randomize whether their conditional or their unconditional contribution is implemented. Through the compensation in the PGG, we incentivize the submission of truthful rank-ordered lists. To guarantee that both the initial submission, as well as the potentially revised preference order are incentive compatible, one of the two is implemented with equal probability to determine the final matching. We incentivize the point beliefs about their partner’s contributions. Participants receive a fixed amount if their stated belief corresponds to the actual unconditional contribution, and they receive no payment otherwise. Additionally, participants are paid based on their performance in the Raven’s matrices task and the loss attitudes elicitation.

Setting and Sample Size The experiment was conducted at the Munich Experimental Laboratory for Economic and Social Sciences (MELESSA). In total, 235 student participants participated in the experiment. The participants were recruited using the online system ORSEE (Greiner, 2015). The experiment was programmed with the software oTree (D. Chen, Schonger, & Wickens, 2016). We conducted 10 sessions (5 sessions per treatment, each with the desired number of 24 participants). On average, participants earned 21.5 EUR (including a show-up fee of 6 EUR), and the experiment lasted around 80 minutes.

2.3 Reciprocal Preferences and the Instability of Matching Markets

Our design reflects the theory that proposers adjust their preferences in *Info* to be matched with a receiver who wants to be matched with them, which in turn leads to a different matching outcome (instability). In pre-registered analyses, we test whether proposers adjust their preferences more often in *Info*, whether these adjustments lead to higher instability, and whether they display reciprocal preferences. Exploratory analyses that were not pre-registered are marked as such.

2.3.1 Instability of the Deferred Acceptance Mechanism

Proposers change their individual preferences more often when they see their potential partners' preferences (*Info*) compared to when they do not see the receivers' preferences (*No-Info*). Regression analysis in Table B.1 in Appendix B.2 confirms that the fraction of preference adjustments is significantly higher in *Info*.

Result 2.1. *The fraction of preference adjustments in Info is 27.67 percent, while it is only 9.67 percent in No-Info. This difference is significant ($p < 0.01$; Mann-Whitney-U test (MWU)).*

As a consequence of more frequent preference adjustments, the fraction of *matching markets* where the rematching outcome changes after the rematching stage is larger under *Info*. Instability of a matching is defined at the *matching market* level. We compute the resulting matching with both the initial and the (potentially) revised preference list. A matching is stable when both resulting matchings are the same (i.e., if all participants are matched to the same partner). Otherwise, a matching is unstable. This implies that a matching market is unstable if at least one of the proposers changed their preferences list, and this change led to a different market outcome. A change in reported preferences

leads to a different outcome only if it results in a proposer-receiver pair that prefers to be matched to each other compared to their current match. The fraction of unstable matchings is substantially larger under *Info* than under *No-Info*¹³ Hence, a *matching market* is nearly four times more often unstable under *Info* than under *No-Info*.

Result 2.2. *There is significantly higher instability in Info than in No-Info. The fraction of unstable matching markets in Info is 40.00 percent; it is only 10.67 percent in No-Info ($p < 0.01$; χ^2 test).*

Thus, we conclude that proposers are more likely to adjust their preference ranking when they see the preferences of receivers, leading to instability in the DA mechanism.

2.3.2 Reciprocal Preferences and Preference Adjustments

Proposer's preference changes indicate the presence of reciprocal preferences. For each preference adjustment, we can classify whether it is consistent with the participants having reciprocal preferences or not. A proposer's preference adjustment is considered *consistent* if the preference order changes, and the (now) more favorably ranked receiver(s) gives a strictly better rank to the proposer compared to the (now) less favorably ranked receiver(s). Formally, this requires that if *Proposer_P* switches the position of *Receiver_R* and *Receiver_S*, and *Receiver_R* was the initially more preferred candidate, then *Proposer_P* must have been ranked strictly better by *Receiver_S* than by *Receiver_R*.¹⁴

Our results strongly support that preference adjustments largely reflect *reciprocal preferences*, and do not reflect proposers' (unsuccessful) attempts to game the strategy-proof mechanism.¹⁵ In *Info*, 73.68 percent of the adjust-

¹³This implies around 24.67 percent of the participants are matched to a new partner in *Info*, while the fraction is 5.67 percent in *No-Info* ($p < 0.001$; MWU).

¹⁴For a formal introduction of reciprocal preferences into matching markets, we refer the interested reader to Opitz and Schwaiger (2022b).

¹⁵This implies that these changes are also not driven by experimenter demand effects or signal confusion of participants on how to integrate the new information into their choices.

ments are consistent with *reciprocal preferences*. This compares to a fraction of 20.69 in *No-Info* where participants could not systematically react to others' preferences.¹⁶ The difference between both conditions is significant ($p < 0.001$; MWU). Table B.2 in Appendix B.2 confirms these findings through a logit regression, documenting a significantly higher likelihood of a consistent preference adjustment (compared to an inconsistent adjustment or none) in *Info*, both in a uni-variate regression (Column 1) and when adding individual-level controls (Column 2).¹⁷

A more detailed exploratory analysis of the determinants of preference adjustments supports the conjecture that proposers' preference adjustments reflect *reciprocal preferences* (see Table B.4 in Appendix B.3). First, the more favorably a proposer ranks their initial partner, the lower the likelihood that the proposer will adjust preferences. This holds true both in *No-Info* (Column 1) and *Info* (Columns 2 & 3). Second, receivers' preferences matter when deciding whether to adjust the preference ranking in *Info*. Being liked by the matched receiver lowers the likelihood that a proposer adjusts their preferences. At the same time, being a preferred candidate by other (non-matched) receivers increases the likelihood of adjusting preferences. Column 2 shows that a more

¹⁶If participants switched the position of two receivers in the preference lists randomly, we would expect 20.9% of the adjustments to be consistent with reciprocal preferences by chance. 24 out of 29 preference adjustments in *No-Info* are such that (only) two receivers switch their position. In more complex cases, the probability of a random adjustment being consistent with reciprocal preferences is even lower.

¹⁷In the loss attitude task (Gächter, Johnson, & Herrmann, 2022) individuals are asked to choose between no payment and a risky lottery with one negative and one positive outcome. Every individual makes several decisions. We keep the positive outcome fixed at 6 Euro, the negative outcomes varies between a loss of 2 and 7 Euro. 2.55 percent of the participants maximize expected payoffs. While the fraction of participants accepting negative expected earnings is negligible (1.28 percent), the vast majority of the participants reject gambles with a positive expected value. The modal response is to accept gambles when the expected value is larger than 2 EUR and reject them otherwise. Loss aversion is defined as the lottery where a participant switches from accepting to rejecting it. For example, if a participant accepts all lotteries, this is coded as 1. If a participant accepts no lottery, this is coded as 7. Cognitive ability is calculated by the number of correctly solved matrices, minus the number of incorrectly solved ones. Out of 10 matrices, participants achieve an average net score of 6.23. 2.55 percent of participants did not solve any matrix correctly, while 5.53 percent solved all 10 matrices correctly.

favorable average rank by the non-matched receivers increases the likelihood of adjusting the preference ranking; Column 3 confirms this pattern by estimating the effect of the best rank received by one of the other three receivers. That proposers in *Info* are less likely to adjust preferences when their matched partner ranked them favorably, and more likely when the other potential partners ranked them favorably is entirely consistent with reciprocal preferences.

Result 2.3. *Preference adjustments are largely reflective of reciprocal preferences in Info, as 73.68 percent of the adjustments are consistent with reciprocal preferences (while this fraction is only 20.69 percent in No-Info).*

Beyond establishing that information about others' preferences leads to higher instability, and that the preference changes are consistent with reciprocal preferences, our design allows us to pin down the underlying reasons for these preference changes.

2.4 Mechanisms Underlying Reciprocal Preferences

In this section, we analyze the reasons underlying reciprocal preferences using a theoretical model, which we then test empirically. In Section 2.4.1, we derive the optimal strategy of a proposer in a stylized version of the experimental *Info* condition and differentiate between belief-based and preference-based mechanisms. In Sections 2.4.2-2.4.4, we put the model's assumptions and implications to the empirical test.

2.4.1 Theoretical Framework

Two proposers (*he*) $p \in \{P, Q\}$ and two receivers (*she*) $r \in \{R, S\}$ participate in a simplified version of the matching market. The DA mechanism forms two teams, each with one proposer and one receiver, to play a PGG.¹⁸ In this model,

¹⁸Section 2.2 offers a detailed description of the PGG and the DA mechanism.

we allow proposers to be altruistic. Each proposer cares about their own direct (monetary) utility $u_p(\pi(c_p, c_r))$ which depends on the monetary payoff $\pi(c_p, c_r)$. The monetary payoff $\pi(c_p, c_r)$ is determined by both partners' contributions $c_{p,r} \in [0, 10]$. Selfish proposers ($a_p = 0$) follow a profit-maximizing strategy and free-ride ($c_p = 0$). Altruistic proposers ($a_p \geq 0$) care not only about their own direct (monetary) utility, but also about their matched partner's direct utility (u_r). The level of altruism $a_p \in [0, 1]$ towards the receiver depends on how likable the proposer perceives the receiver to be.

The core of our experimental treatment *Info* is that applicants learn how receivers rank them. We make two main assumptions about why this matters. First, we assume that the level of altruism is determined by the proposer's initial assessment of the receiver (l_p), and on how likable the receiver perceives him (l_r) to be. A higher level of l_p and l_r increases the level of altruism.¹⁹ In other words, we assume that agents are more altruistic towards partners they like (Leider et al., 2009) and that “receiving information that another is attracted to you is a powerful determinant of liking” (Montoya & Horton, 2012). In our context, we assume that the receiver's rank is informative about l_r .²⁰

Assumption 2.1. *Preference-based mechanism: The level of altruism (a_p) increases in l_r .*

Second, we assume that a proposer expects a receiver who ranked him favorably to contribute more to the PGG.²¹

Assumption 2.2. *Belief-based mechanism: The belief about a receiver's contributions (\hat{c}_r) increase in l_r .*

The direct (monetary) utility function $u_{p,r}$ is positive, monotonically increasing, continuous, and concave in the monetary payoff $\pi_{p,r}$ and has the same

¹⁹We assume that $l_{r,p}$ is a natural number.

²⁰This is somewhat related to the idea of R. Dur (2009) that agent i 's altruism towards another agent j depends the altruism of agent j towards agent i (which agent i infers from some action of agent j).

²¹For simplicity, we assume that the proposer knows the contribution c_r after observing l_r .

functional form for all agents. The adjusted utility of a proposer is given by:²²

$$v_p = u_p(\pi_p(c_p, c_r)) + a_p(l_p, l_r) \cdot u_r(\pi_r(c_p, c_r))$$

The model illustrates 1) how a proposer optimally selects his partner and 2) how he decides about his contributions to the PGG. The timing of the model mirrors our experimental design in *Info*. First, proposers and receivers submit their preferences to a mechanism. At this point, the proposer has no information about l_r , his belief is the same for both receivers ($\hat{l}_R = \hat{l}_S$). This implies that proposers base their decision solely on l_p . Then, proposers learn the true preferences of both receivers (l_R, l_S). As proposers have (a priori) no information about l_r , a first-place ranking provides a weakly positive update about l_r while a second-place ranking presents a weakly negative update about l_r . Afterwards, proposers can adjust their ranking and play the PGG with their matched receiver. We solve the model by backward induction, first describing the contribution decisions before examining the implications for preference changes.

When matched with a receiver, a proposer optimizes by choosing his contribution to the PGG. Increasing the contribution level lowers his monetary outcome while raising the matched receiver's payoff. The proposer's adjusted utility is maximized if the decrease in his marginal direct utility equals the increase in the matched receiver's marginal utility times the altruism factor towards her (altruism utility).

$$\max_{c_p} v_p : u_p(\pi_p(c_p, c_r)) + a_p(l_p, l_r) \cdot u_r(\pi_r(c_p, c_r)) \quad (2.1)$$

$$\frac{\partial v_p}{\partial c_p} = \underbrace{\frac{\partial u_p}{\partial c_p}}_{<0} + \underbrace{a_p(l_p, l_r) \cdot \frac{\partial u_r}{\partial c_p}}_{>0} = 0 \quad (2.2)$$

²²The idea of direct (monetary) utility and adjusted utility is first described by Levine (1998).

Following the optimization problem of the proposer²³, we give a short overview of the model's main propositions. These proofs can be found in Appendix B.1.

Proposition 2.1. *An increase in l_r has a non-negative effect on the contribution of a proposer c_p .*

We assume that the level of altruism a_p (Assumption 2.1) and belief about the receiver's contribution \hat{c}_r (Assumption 2.2) increase in l_r . If l_r increases, both channels then increase the proposer's contribution c_p given an interior solution. First, as the level of altruism a_p increases, a proposer benefits more from the receiver's monetary payoff. Hence, the proposer's contribution c_p increases. Second, the higher contribution of the receiver decrease the proposer's marginal direct (monetary) utility and increase the receiver's marginal direct monetary utility. To equalize these marginal benefits (weighted by the altruism factor), the proposer increases his contribution.

Proposition 2.2. *A change of preferences for proposer P can only happen if a receiver R , whom proposer P initially ranked worse than receiver S , ranks him better than receiver S .*

If the proposer observes that he is ranked first by a receiver, he positively updates l_r . This change increases the proposer's adjusted utility of being matched with the receiver. Through a higher l_r (and hence a higher a_p and \hat{c}_r), the proposer both expects a higher monetary outcome for himself and cares more about the receiver. Both effects result in a higher contribution and lead to a higher utility for the proposer.

We can now derive the proposer's preferences over receivers and show why these preferences may change. A proposer ranks receivers based on his expected adjusted utility v_p of being matched with them if a strategy-proof mechanism is applied. His preference order can change upon learning how the receivers rank him. A positive update about l_r (weakly) increases the adjusted utility of being

²³The second order condition holds ($\frac{\partial^2 v_p}{\partial c_p^2} = \frac{\partial^2 u_p}{\partial c_p^2} + a_p \frac{\partial^2 u_r}{\partial c_p^2} < 0$).

matched with a receiver. The reverse is true for a negative update. Therefore, a change of preferences can, for example, happen if the proposer initially ranked receiver R over receiver S , but then learns that he was ranked first by receiver S , and second by receiver R . This can, but need not, change the proposer's preference order. For an altruistic proposer ($a > 0$), these changes can be driven by preference-based and belief-based motives. For selfish proposers ($a = 0$), changes are entirely driven by beliefs about others' contributions. Selfish proposers will never contribute, but want to be matched to the highest contributing receiver.

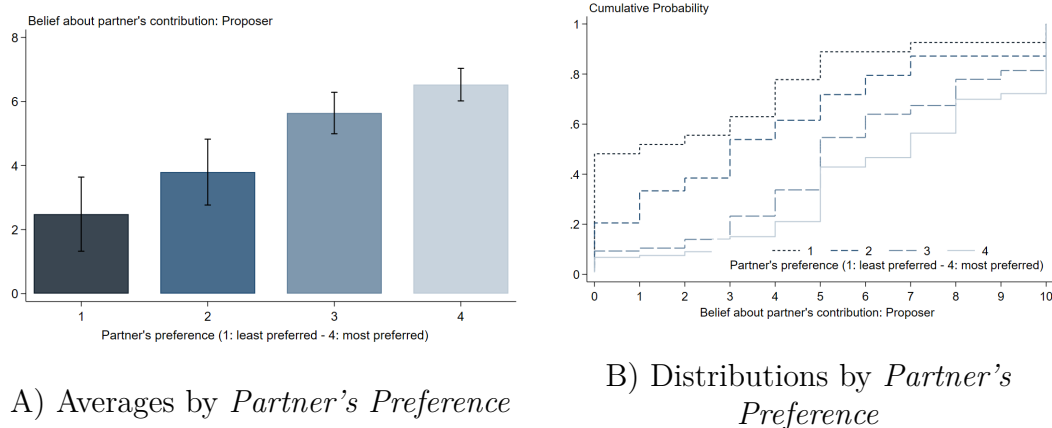
Our model predicts preference changes consistent with reciprocal preferences (see Results 2.1 and 2.3). It highlights two channels for this behavior. First, participants change their preferences because they expect partners who like them to contribute more to the PGG. Preferences are interpreted as a signal about the match-specific value, and proposers change their preferences accordingly (belief-based). Second, a proposer may prefer to be matched with a receiver who liked them because he is more altruistic towards such a receiver (preference-based). Our results on the PGG behavior allow us to test whether preference-based or belief-based reasons explain the adjustments in *Info*, and how these adjustments translate into cooperative behavior.

2.4.2 Evidence for a Belief-Based Mechanism

We test the belief-based channel by analyzing (incentivized) beliefs of proposers about their matched receivers' contributions depending on how their partner ranked them. This means that we directly test our model's key Assumption 2.2 – that the receivers' preferences (l_r) are perceived as a signal about their contributions (\hat{c}_r). We first show that the receiver's preferences are indeed perceived as a signal about their contribution. We then demonstrate that these beliefs are accurate.

We find that proposers expect receivers who rank them better to contribute more to the PGG. Figure 2.2 shows this plotting beliefs over *Partner's prefer-*

ences (1-4). This variable takes the value of four if the proposer was the matched receiver’s most preferred choice, three if the participant was the second most preferred choice, and so on. Panel A shows that mean beliefs about the matched receiver’s contribution increase with the receiver’s preferences. Panel B illustrates this trend by presenting cumulative distribution functions. It shows, for example, that only 6.77 percent of proposers believe that their partner will contribute nothing when they were their partner’s first choice. By comparison, 48.15 percent believed their partner will not contribute anything to the public good when they were their partner’s least preferred choice.



Notes. This figure displays the beliefs of proposers in *Info* about the unconditional PGG contributions of their matched receiver by the preferences of the matched partner. *Partner's preferences (1-4)* takes the value of four if the participant was the most preferred choice of their matched partner, three if the participant was the second most preferred choice, and so on. Panel a) shows averages, Panel B) the cumulative distribution functions.

Figure 2.2: Beliefs about Receiver’s PGG Contributions: Proposers in *Info*

Table 2.1 corroborates that proposers expect receivers who like to be matched with them to contribute more. The effects are sizable (Column 1), and remain so when controlling for the round and individual-level characteristics of the proposer (Column 2). Proposers expect matched receivers to contribute around 1.4 Taler (out of 10) more if they are ranked one place better on the receiver’s preference list. This expectation is consistent with the notion that the expression of interest is “one cue to identify someone who is likely to act [...] cooperatively” (Montoya & Insko, 2008, p.478). Given such beliefs, a proposer may expect a

change in their preference order to be payoff-maximizing if it results in being matched with a receiver who prefers them as a partner.

Table 2.1: PGG Behavior of Proposers in *Info*

	Belief Partner Contribution		Unconditional PGG Contribution		Avg. Conditional PGG Contribution	
	(1)	(2)	(3)	(4)	(5)	(6)
Partner's preference (1-4)	1.348*** [.977,1.720]	1.382*** [.915,1.849]	.771*** [.350,1.193]	.794*** [.340,1.248]	.415*** [.128,.702]	.416*** [.147,.685]
Preference for partner (1-4)	-.073 [-.437,.291]	-.059 [-.445,.328]	.105 [-.219,.429]	.146 [-.172,.463]	.013 [-.182,.209]	.026 [-.159,.212]
Round		.064 [-.118,.245]		-.197*** [-.338,-.055]		-.119** [-.214,-.023]
Loss Aversion		-.795*** [-1.389,-.201]		-.710** [-1.408,-.012]		.040 [-.429,.509]
Cognitive Ability (Raven's)		.500* [-.037,1.037]		.337 [-.347,1.021]		-.159 [-.557,-.239]
Male		-.025 [-1.434,1.383]		-.589 [-2.273,1.096]		-.605 [-1.581,-.371]
Observations	285	285	285	285	285	285

Notes. OLS Regressions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals.

Partner's preferences (1-4) takes the value of four if the participant was the most preferred choice of their matched partner, three if the participant was the second most preferred choice, and so on. *Preference for partner (1-4)* takes the value of four if the matched partner was the first choice of the participant, three if the matched partner was the second choice, and so on. *Round* is a count variable, indicating the number of the current round (Round 1-5). *Loss aversion* and *Cognitive ability* are calculated as detailed in Footnote 17. *Male* is an indicator taking the value of 1 if a participant indicated to identify as male.

Result 2.4. *Proposers expect a significantly higher unconditional contribution from receivers who rank them better ($p < 0.01$).*

The beliefs of proposers are in line with receivers' actual cooperation behavior. Figure B.2 (Appendix) displays that receivers contribute more to the PGG when matched with proposers they prefer. Table B.3 (Appendix) shows that each rank the matched proposer is up in the preference list leads to a 0.96 Taler higher contribution to the PGG in the preregistered specification of Column 2, Table B.3.²⁴ Thus, proposers correctly expect receivers' preferences to influence their contribution decisions.

²⁴Our design does not allow us to disentangle the underlying reasons for higher contributions by receivers. Still, our data is consistent with receivers (partly) contributing more when they like their partner, because they expect proposers they rank favorably list to contribute more to the PGG than proposers they rank less favorably (see Figure B.3 and Table B.3 in the Appendix). As receivers do not know that proposers learn their given rank, receivers believe that they can identify high-contributing proposers. In light of our findings that none of the personality questions predicts contributions in the PGG (see Table B.5), these beliefs turn out to be wrong.


In sum, we provide evidence for a belief-based mechanism underlying reciprocal preferences. We show that proposers rationally expect higher contributions from receivers who rank them favorably, which provides a rationale for the observed preferences adjustments.

2.4.3 Evidence for a Preference-Based Mechanism

We test whether preference-based explanations play an additional role for reciprocal preferences by analyzing proposers' conditional contributions. Conditional contributions are independent from beliefs about the partners' contribution and, therefore, directly informative about the level of altruism (a_p). If proposers conditionally contribute more when interacting with a receiver who ranked them favorably, this implies higher altruism. Hence, we can directly test whether altruism is sensitive to the partner's preferences (l_r), which we presume by Assumption 2.1.

Proposers provide higher conditional contributions when matched to a receiver who ranks them favorably. Their average conditional contributions increase monotonically in the position on the receiver's preference list (see Figure 2.3, Panel A). Across the eleven conditional contribution decisions, they contribute around 0.4 Taler more for each spot they are ranked better (see Table 2.1). These averages mask an interesting heterogeneity, which we investigate in an exploratory analysis. Figure 2.3, Panel B shows that this difference in behavior is especially pronounced when facing higher contributions of the partner. The sub-figure plots the regression coefficient of the partner's preferences for each of the eleven contribution decisions, given the specification in Table 2.1, Column (6). For low contribution values of the receiver, the receiver's preferences do not strongly impact proposers' behavior. For example, proposers do not condition their contributions on whether a free-rider wants to be matched with them or not. However, receivers' preferences become an important determinant of proposers' conditional contributions when receivers make higher contributions. Proposers are then more altruistic towards receivers that indicate

a preference to be matched with them.



Notes. This figure displays the average conditional contributions of proposers in *Info* by the preferences of matched receiver.

Notes. The figure plots the regression coefficients β_1 of the regressions $y_i = \beta_1 * Partner's\ preference + \beta_2 * Preference\ for\ partner + \beta_3 * t + \beta_4 * X_p$, corresponding to Table 2.1 with t indicating the round, and X_p as a vector consisting of gender, cognitive ability, and loss aversion. The outcome variables y_i is the conditional contributions of a proposer for any (unconditional) contribution $i \in 0, 10$ of the matched receiver. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

A) Averages by *Preference for Partner*

B) Coefficient Plot

Figure 2.3: Average Conditional PGG Contributions: Proposer in *Info*

Result 2.5. *The conditional contributions of proposers are significantly higher when they interact with a receiver who ranked them favorably, especially for high levels of contributions by the receiver ($p < 0.01$).*

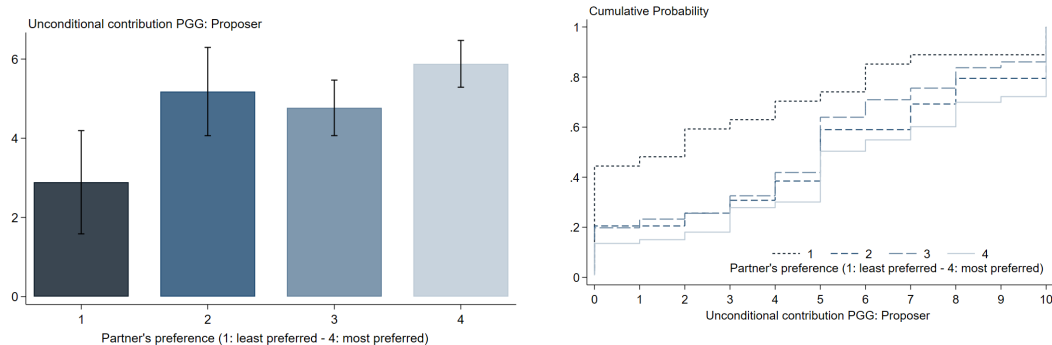
This provides evidence that preference-based explanations are important for the observed behavior. We document higher social preferences towards receivers who rank the proposer favorably. Reciprocal preferences are therefore likely to stem both from preference-based and belief-based factors.

2.4.4 Unconditional Cooperation

Unconditional contribution decisions inform about the overall effect of reciprocal preferences in one-time simultaneous cooperation. Higher altruism (preference-based) leads to higher unconditional contributions. Higher beliefs about the

contributions of the partner (belief-based) result in higher unconditional contributions by those willing to contribute more the more the other contributes (*i.e.*, *conditional cooperation* as in Fischbacher, Gächter, & Fehr, 2001). The analysis of unconditional contributions directly tests Proposition 2.1.

On average, proposers contribute more to the PGG when interacting with a receiver who ranks them favorably. Table 2.1, Column 4 documents that proposers contribute around 0.8 Taler more when they are ranked one spot more favorably by their matched receiver. The partner’s preferences (see l_r in the model) are more predictive of the actual contribution behavior of proposers than their own (initial) preference for the partner (l_{pr}). Figure 2.4 shows that unconditional contributions are especially low when interacting with a receiver who ranked them on the worst spot of their preference list.²⁵



A) Averages by *Partner's preference*

B) Distributions by *Partner's preference*

Notes. This figure displays the unconditional contributions of proposers in *Info* by the preferences of the matched receiver. *Partner's preferences (1-4)* takes the value of four if the participant was the most preferred choice of their matched partner, three if the participant was the second most preferred choice, and so on. Panel a) shows averages, Panel b) the cumulative distribution functions.

Figure 2.4: Unconditional PGG Contributions: Proposers in *Info*

Result 2.6. *The unconditional contribution of proposers is significantly higher when they interact with a receiver who ranked them favorably on their preference list ($p < 0.01$).*

²⁵This is consistent with evidence from other domains which highlights the aversion of being ranked last, such as Kuziemko et al. (2014).

Comparing social efficiency between treatments, we find that average (unconditional) cooperation and payoffs are higher in *Info* than in *No-Info*. While proposers in *No-Info* contribute on average 4.12 (out of 10) Taler to the PGG, contributions are around 25.7 percent higher in *Info* (5.18). On average, participants in *Info* contribute 0.96 Taler more to the PGG ($p = 0.039$; MWU), which translates into 0.48 Taler higher payoffs in *Info*. Accordingly, information about others' preferences increases average cooperation and payoffs.²⁶

2.4.5 Mechanisms across Treatments

In the previous analyses of the mechanisms underlying reciprocal preferences, we compared proposers ranked favorably to proposers ranked less favorably by their partner within *Info*. To corroborate these results and to substantiate that they are specific to the information environment in *Info*, we hold the proposer's rank received across treatments constant. We can compare beliefs and contributions in the situation in which proposers knew their partner's preference (*Info*) to that in which the proposers did not know it (*No-Info*). Hence, in a sort of Placebo test, we estimate the effect of *knowing* the rank on contributions and beliefs. Table 2.2 shows that none of our variables of interest is significant in *No-Info*.

²⁶Figure B.4 (Appendix) shows that the higher average payoff does not mask a substantial mean-variance trade-off. The treatment *Info* increases payoffs across the distribution.

Table 2.2: PGG Behavior of Proposers in *Info* and *No-Info*

	Belief Partner Contribution	Unconditional Contribution	Avg. Conditional Contribution
	(1)	(2)	(3)
Preference for partner (1-4)	.246 [-.059,.551]	.129 [-.172,.430]	.035 [-.128,.198]
Partner's preference (1-4)	.106 [.265,.477]	-.068 [-.499,.363]	-.053 [-.339,.233]
Partner's Preference X Info	1.208*** [.600,1.815]	.862*** [.248,1.475]	.444** [.062,.826]
Info	-2.566** [-4.704,-.429]	-1.609 [-3.826,.609]	-.763 [-2.123,.596]
Round	-.021 [-.141,.099]	-.254*** [-.359,-.149]	-.169*** [-.238,-.100]
Loss Aversion	-.609** [-1.112,-.106]	-.570** [-1.130,-.009]	.300* [.041,.641]
Cognitive Ability (Raven's)	.272 [-.133,.678]	.180 [-.315,.675]	-.196 [-.465,.073]
Male	-.928* [-1.974,.117]	-.884 [-2.155,.387]	-.325 [-1.118,.468]
Observations	575	575	575

Notes. OLS Regressions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals.

Partner's preferences (1-4) takes the value of four if the participant was the most preferred choice of their matched partner, three if the participant was the second most preferred choice, and so on. *Preference for partner (1-4)* takes the value of four if the matched partner was the first choice of the participant, three if the matched partner was the second choice, and so on. The interaction term *Partner's Preference X Info* takes the value of zero for observation in *No-Info*, and the value of *Partner's Preference X (1-4)* in *Info*. *Info* is an indicator, taking the value of one if the participant was randomly assigned to the treatment *Info*. *Round* is a count variable, indicating the number of the current round (Round 1-5). *Loss aversion* and *Cognitive ability* are calculated as detailed in Footnote 17, *Male* is an indicator taking the value of 1 if a participant indicated to identify as male.

2.4.6 Gender Heterogeneity

Overall, male participants drive the differences in proposer's behavior depending on their partner's preference. In an exploratory regression analysis (Table 2.3), we show that men's contribution decisions are significantly more influenced by their partner's preferences. This is true for both their unconditional (Column 2), and their conditional contributions (Column 3). In addition, men's beliefs (Column 1) about others' contributions are more responsive to their position on their partner's preference list.²⁷

This gender heterogeneity raises interesting questions regarding how men and women react to rankings and evaluations. Previous research has found that women update more pessimistically than men when receiving negative feedback (Berlin & Dargnies, 2016). In addition, women attribute negative feedback

²⁷This leads to a lower average belief accuracy for men than for women (average deviation of 3.54 vs. 4.47, $p = 0.02$; MWU).

to skill rather than to luck more often than men (Shastry, Shurchkov, & Xia, 2020), and react more strongly to likeability ratings based on their appearance (Gerhards & Kosfeld, 2020). On the other hand, our results suggest that men take the ranking more “personally” and react to it more strongly. This is consistent with previous findings recognizing women as being more ego-defensive (Möbius et al., 2022), and as having stronger internalized norms about giving, which leads to a lower elasticity of their altruism (Andreoni & Vesterlund, 2001). It is also in line with the finding of Barankay (2012) that feedback about performance rankings changes the behavior of men, but not of women.

Table 2.3: Gender Heterogeneity of Proposers in *Info*

	Belief Partner Contribution	Unconditional Contribution	Avg. Conditional Contribution
	(1)	(2)	(3)
Preference for partner (1-4)	-.086 [-.462,.289]	.107 [-.225,.438]	.006 [-.174,.186]
Partner’s preference (1-4)	1.004*** [.457,1.551]	.253 [-.225,.731]	.134 [-.106,.373]
Partner’s preference X Male	.798* [-.037,1.634]	1.142*** [.360,1.925]	.595** [.123,1.068]
Male	-2.568* [-5.553,.418]	-4.225*** [-7.246,-1.205]	-2.500** [-4.478,-.522]
Controls [Round + Individual]	Yes	Yes	Yes
Observations	285	285	285

Notes. OLS Regressions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals.

Preference for partner (1-4) takes the value of four if the matched partner was the first choice of the participant, three if the matched partner was the second choice, and so on. *Partner’s preferences (1-4)* takes the value of four if the participant was the most preferred choice of their matched partner, three if the participant was the second most preferred choice, and so on. *Male* is an indicator taking the value of 1 if a participant indicated to identify as male. *Partner’s preferences x Male* takes the value of zero for observation with $Male=0$, and the value of *Partner’s Preference X (1-4)* when $Male=1$. All columns control for *Round*, which is a count variable, indicating the number of the current round (Round 1-5), as well as *Loss aversion* and *Cognitive ability* that are calculated as detailed in Footnote 17.

2.4.7 Similarity, Homophily, and Reciprocal Preferences

Perceived similarity influences behavior in various decisions (e.g., Y. Chen & Li, 2009; Eckel & Grossman, 2005; Hedegaard & Tyran, 2018), and has been shown to relate to interpersonal attraction (McWhirter & Jecker, 1967). Hence, the effect of partner’s preferences’ on their behavior may operate through a channel of homophily (Currarini, Jackson, & Pin, 2009; McPherson, Smith-Lovin, & Cook, 2001). If a proposer only has an imprecise signal about their similarity with the matched receiver based on five questions, the receiver’s preferences

(that are based on five different questions) may provide a signal about their similarity. Assuming common preferences to interact with a similar individual, the preference of the partner can be interpreted as information about their similarity.²⁸ So far, we have shown that information about others' preferences leads to more instability in matching markets through preferences adjustments (Section 2.3.1), that these adjustments are consistent with reciprocal preferences (Section 2.3.2), and that these adjustments likely stem from a combination of belief-based and preference-based factors (Sections 2.4.2 & 2.4.3). However, we have not yet established that the partner's preferences are not only similarity signals, but a fundamental determinant of behavior.

In Table 2.4, we provide evidence that partner preferences matter beyond being a signal for similarity. We calculate similarity as the inverse of the average distance between the questionnaire responses of the matched partners (*Manhattan distance*). For example, the value is equal to 0 if one of the partners clearly affirmed each statement and the other clearly rejected all (i.e., the difference of their answers on the four-point Likert scale is maximal), and it is equal to 3 if they answered each question identically. First, the main coefficient of the partner's preferences remains constant when controlling for similarity based on all 15 questionnaire answers (Column 2). This implies that the partner's preferences do not fully operate by providing an accurate signal regarding similarity. Second, the main coefficient remains constant when including the similarity measure based on the five randomly selected questions for which the proposer has seen their partner's responses (Column 3). The positive and significant similarity coefficient implies that proposers condition their contributions on whether their partner's responses match their own. At the same time, the similar main coefficients in Columns 2 and 3 imply that there is little additional signaling value in the preferences of the other agent. If there were, we would expect the main coefficient in Column 3 to be substantially higher than in Col-

²⁸Similar to Currarini and Mengel (2016), we find that similarity is an important predictor for partner choice in the PGG. The raw correlation between the rank given to a receiver and our baseline measure for dissimilarity (Manhattan Distance) is 0.23, $p < 0.001$.

umn 2. Third, the coefficient remains stable when we control for the similarity in answers across the five randomly selected questions to which the receiver has seen the proposer’s answers. If preference were a signal about this similarity, this would again imply a lower main coefficient in Column 4 than in Column 1 or 2.

Table 2.4: Homophily and Unconditional Contributions of Proposers in *Info*

	Unconditional PGG Contribution (0-10)			
	(1)	(2)	(3)	(4)
Partner’s preference (1-4)	.794*** [.340,1.248]	.704*** [.252,1.155]	.812*** [.371,1.254]	.794*** [.344,1.245]
Preference for partner (1-4)	.146 [-.172,.463]	.044 [-.313,.402]	.123 [-.195,.441]	.139 [-.181,.458]
Similarity Answers (0-3) [Manhattan]		1.357 [-.563,3.278]		
Similarity of Shown Answers (0-3) [Manhattan]			1.041** [.115,1.968]	
Similarity of Receiver’s Answers (0-3) [Manhattan]				.862 [-.293,2.017]
Controls [Round + Individual]	Yes	Yes	Yes	Yes
Observations	285	285	285	285

Notes. OLS Regressions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals.

Partner’s preferences (1-4) takes the value of four if the participant was the most preferred choice of their matched partner, three if the participant was the second most preferred choice, and so on. *Preference for partner (1-4)* takes the value of four if the matched partner was the first choice of the participant, three if the matched partner was the second choice, and so on. *Similarity* is the inverse of the average distance between the questionnaire responses of the matched partners (Manhattan distance). *Similarity Answers (0-3)* is calculated based on all 15 questionnaire items, *Similarity of Shown Answers (0-3)* is based on the five questions the proposer saw the partner’s answers for, *Similarity of Receiver’s Answers (0-3)* is based on the five questions the receiver saw the partner’s answers for. All columns control for *Round*, which is a count variable, indicating the number of the current round (Round 1-5), the indicator *Male* taking the value of 1 if a participant indicated to identify as male, as well as *Loss aversion* and *Cognitive ability* that are calculated as detailed in Footnote 17.

2.5 Discussion

This paper shows that reciprocal preferences represent a powerful source of instability in matching markets. First, we demonstrate that reciprocal preferences exist – that is, that participants *like to be liked*. When participants learn the preferences of their potential partners, they adjust their preferences and rank more favorably those who would like to be matched with them. Second, we show that these changes substantially increase the number of unstable matchings. Third, we investigate the underlying motives of reciprocal preferences and

find evidence for both belief-based and preference-based mechanisms. On the one hand, proposers expect receivers who like them to contribute more to the PGG. This provides a profit-maximizing rationale for preference adjustment due to changes in beliefs. On the other hand, proposers are more altruistic towards receivers who like to be matched with them. This supports a preference-based rationale for reciprocal preferences.

The PGG reflects the cooperative nature of many matching markets. In matching markets, not only are relationships formed without the coordinating function of prices, but also within the relationships there are non-contractible elements. Insofar as these elements relate to effort provision and commitment, cooperation plays a crucial role in these relationships. Consider a university that wants to hire an enthusiastic job market candidate, and (in turn) a candidate who also wishes to receive support from the department. Both choosing a cooperative partner *and* being in a relationship where one wants to be cooperative oneself is key in such a setting, where decisive aspects cannot be contracted upon. The PGG allows us to investigate both of these channels.

Notwithstanding, our stylized experimental setting does not capture all aspects of the preference-based foundation of reciprocal preferences. Real interactions put more weight on psychological mechanisms, such as the non-pecuniary disutility of working with someone who does not like you. Therefore, investigating reciprocal preferences in inter-personal coordination tasks constitutes an avenue for future research. Compared to our experimental design, which likely provides a lower bound for the effect of reciprocal preferences, the effects could be even more pronounced when individuals expect a personal interaction.

Our results are policy-relevant, as they contribute to a better understanding of matching markets, cooperative behavior, and effective team formation in organizations. First, our results can help to design matching markets more efficiently. It is necessary to understand why matching markets sometimes fail to reach their full potential. Opitz and Schwaiger (2022b) theoretically show that reciprocal preferences can be a source of instability. Evidence from real-world

matching markets suggests that reciprocal preferences play an important role. Nevertheless, observational data does not allow for teasing apart reciprocal preferences, uncertainty, and other potential reasons for market failures. This paper establishes the empirical relevance of reciprocal preferences and thus highlights the importance of information design in matching markets. While learning about others' real-world preferences might sometimes be more subtle than in the experiment, already observing the final matching can lead to updates about other participants' preferences and result in instability. Our sizeable effects suggest that reciprocal preferences also play an essential role in slightly different information environments. Understanding the importance of reciprocal preferences helps to reconcile strategic modifications of the theoretically efficient mechanism by participants (e.g., by offering early admission or making admission decisions contingent on others' preferences). In addition, it helps to design mechanisms that accommodate agents' reciprocal preferences.

Second, we enhance understandings of social preferences and social proximity. Previous (Leider et al., 2009) has established that we treat those close to us more favorably, without being able to differentiate between our liking, being liked, and similarity. By isolating the role of being liked, we provide evidence that giving in a relationship depends on not only our own preferences, but also others' preferences. These findings are consistent with literature outside economics that emphasizes the wish to be liked as a universal desire (Baumeister & Leary, 1995) with neural underpinnings (Davey et al., 2010), and the susceptibility of our own interpersonal preferences to the preferences of others (Montoya & Horton, 2012, 2014). We demonstrate that this susceptibility implies that interpersonal preferences are another currency of reciprocity, expanding previous findings on which type of gifts can lead to productivity gains (e.g., Kube, Maréchal, & Puppe, 2012). Hence, we link interpersonal preferences to organizational implications for motivating workers.

Third, our findings on the relevance of reciprocal preferences have broader organizational implications for team formation and teamwork. Organizational

processes and production steps require voluntary cooperation to achieve optimal results (Deversi, Kocher, & Schwieren, 2020). We show that being liked can be necessary for cooperation. Previous literature has established that self-selected teams display homophily in their traits and networks, leading to higher satisfaction and effort (Boss et al., 2021; R. Chen & Gong, 2018). We provide a foundation for these results by highlighting greater cooperation when collaborating with a partner who likes you. We show that even in a stylized setting without personal interactions, we observe homophily in sorting, higher cooperation among those matched with partners who like them, and higher profits.

Chapter 3

Reciprocal Preferences in Matching Markets¹

3.1 Introduction

Standard matching theory assumes that agents do not care about the preferences of their potential partners. In this paper, we study the observation that individuals *like to be liked*. For example, school principals “want to run a school where people want to be there”, and hence “take into account [...] that one kid wants to go there more than another kid”.² The same holds true in the labor market—employees prefer to work for firms that favor them, and may even reconsider a job offer after learning that they were not the first-choice candidate.³ Conversely, an employer may look for a worker who prefers to work for them rather than for another company. Agents who prefer to be matched to a partner who likes to be matched with them are defined as having *reciprocal preferences*.

In this paper, we introduce reciprocal preferences, apply them to match-

¹This chapter is based on joint work with Timm Opitz.

²David M. Herszenhorn, Council Members See Flaws in School-Admissions Plan New York Times, Nov. 19, 2004, <https://www.nytimes.com/2004/11/19/education/council-members-see-flaws-in-schooladmissions-plan.html>, accessed 01/31/2022).

³See <https://www.forbes.com/sites/lizryan/2018/01/20/im-the-second-choice-candidate-should-i-still-take-the-job>, accessed 01/31/2022, for an example on the perceived importance of being the most preferred candidate.

ing markets, and analyze implications for mechanism design. We resort to the standard marriage model (Gale & Shapley, 1962), where agents are one-to-one matched. We augment the setting by allowing reciprocal preferences on one side of the market. These agents care about the preferences of the agents on the other side but they do not know these preferences perfectly. We analyze both standard two-sided markets and school choice settings, in which one side of the market is non-strategic. In both settings, we investigate the implications of reciprocal preferences on stability when the Deferred Acceptance (DA) mechanism is applied (Gale & Shapley, 1962). The DA mechanism plays a key role in two-sided matching markets because it achieves Pareto efficiency under standard assumptions, such that no participant benefits from breaking up the formed match. Moreover, we generalize and analyze stability in a broad class of matching mechanisms.

We derive three main results. First, stability of the DA mechanism in two-sided matching markets ceases to hold when agents care about others' preferences over themselves without perfectly knowing these preferences. The DA mechanism achieves stability in a standard setting under complete information. Under incomplete information, preference misrepresentations of agents may lead to instability. In our setting with incomplete information and reciprocal preferences, instability arises through an updating about the preferences of other agents in the market. However, this instability is not due to strategic play but is due to the disclosure of the other side's preferences. We test stability under two different degrees of information revelation. Under the more restrictive notion, agents only observe the final matches. Under the less restrictive notion, they also learn their matched partner's type. Under both stability notions, agents might infer the true preferences (and thus the types) of other agents, which causes instability.

Second, we show that there is no alternative mechanism that guarantees a stable matching in two-sided markets with (one-sided) reciprocal preferences and uncertainty about true preferences. Stability cannot be achieved because

agents update about each others' preferences by seeing the final matching. We show that this also holds when agents additionally learn the preferences of their matched partner.

Third, modified versions of the DA mechanism achieve stability in a school choice setting where one side of the market must state its true preferences based on laws and regulations, which is the case with schools. We show that the standard DA mechanism with reciprocal preferences does not prevent instability in school choice settings. The students still face uncertainty about the schools' preferences and can only infer them after the matching. Alternative mechanisms can resolve the problem of uncertainty, which can be either a sequential variant of the DA mechanism where students learn schools preferences or a variant of the DA mechanism that allows students to state their complete reciprocal preference profile.

These results help us to understand why certain matching markets do not work satisfactorily. If reciprocal preferences are strong and standard matching mechanisms do not consider these, then involved parties may be reluctant to adopt a centralized matching mechanism. They may instead prefer decentralized markets because they allow them to learn the preferences of others.⁴ Agents with reciprocal preferences may also have incentives to modify an existing matching mechanism. To attract especially interested candidates, universities introduce early admissions (Avery & Levin, 2010), give a bonus to students who rank the institution well, or even only accept these candidates (Chiu & Weng, 2009).⁵ While individually rational, these modifications may prevent the desired functioning of matching mechanisms. If a mechanism does not achieve the goal of finding the best possible partner and modifications are not feasible, then it may even fail completely after some time (McKinney, Niederle, & Roth,

⁴For example, Gundlach (2021) documents that only three percent of parents of kindergarten-age children in Germany approve the use of an algorithm alone when deciding on the allocation of daycare places.

⁵The idea that less interested candidates will receive a deduction has been formally incorporated in the centralized mechanism for high school admission in Taiwan (U. Dur et al., 2022).

2005).

We derive important policy implications for the efficient design of markets when agents have reciprocal preferences by combining our results of two-sided mechanisms and the school choice setting. The feasibility to achieve stable allocations in a school choice setting implies that it may be advantageous to move a standard two-sided market closer to a school choice setting. Students in a school choice setting can be helped if the (binding) admission criteria are known before the mechanism takes place and if they have information about the past success criteria. Given that schools are non-strategic, this eliminates uncertainty about their true preferences and helps the other market side to submit their ranking. In centralized matching mechanisms where one side is an institution connected to the market designer, it can be feasible to force one side of the market to state their true preferences. Additionally, we point to interesting trade-offs in the information design of matching mechanisms. The communication of the final matching can cause allocations to break apart that would otherwise be stable without this information.

Our general framework assumes that the preferences of potential partners are associated with their type, which is relevant to an agent. By considering reciprocal preferences, we are able to focus on situations in which an agent cares about being liked by their partner. However, the framework applies to all situations in which types of agents differ in their preferences and other agents care about these types. This does not even require agents who care about preferences over themselves. For example, two preference profiles that rank an agent the same but differ in how they rank other agents may signal an underlying type that the agent cares about.

Our findings connect to two strands of the literature of matching markets. First, we contribute to an emerging literature of non-standard preferences in matching markets by introducing reciprocal preferences. Fernandez (2020) highlights that regret aversion may induce truth-telling for both market sides in the DA mechanism. Other studies emphasize how non-standard preferences pre-

vent the desired functioning of allocation mechanisms. These studies show that costly information acquisition about one's preferences leads to higher acceptance rates of early offers, despite not being more desirable (Grenet, He, & Kübler, 2022), and that expectation-based loss aversion can lead to non-truthful preference submissions (Dreyfuss, Heffetz, & Rabin, 2021; Meisner & von Wangenheim, 2021). Meanwhile, Antler (2015) finds that a slight modification of the standard DA mechanism preserves stability when the agents' preferences are directly affected by the reported preferences of others. In his model, agents have perfect information about others' preferences, while outsider observers do not. The outside observers only see the submitted rankings. Because agents have image concerns, they prefer to be matched with a partner who ranked them highly. This implies that even if an agent knows that their counterpart does not prefer to be matched to them, they do not care as long as they are highly ranked. In contrast, our analysis assumes that agents care about the unknown *true* preferences and the types of their potential partners, rather than the stated preferences in an environment of complete information. In complementary work (Opitz & Schwaiger, 2022a), we validate the assumption of reciprocal preferences experimentally, and show that agents indeed care about the preferences of their partners.

Second, we contribute to the literature of incomplete information in matching markets by studying an environment without perfect knowledge about the preferences of the potential partners. Roth (1989) introduces uncertainty about the preferences of other players. However, the preferences of an agent are not influenced by other players' preferences. Roth (1989) instead analyzes stability under the assumption that all preferences were to become common knowledge and finds that no mechanism is stable with respect to the true preferences. The instability arises due to strategic play. If an agent has enough information about the preferences of other players, then stating non-truthful preferences can be optimal. An agent misrepresents their preferences to reject a candidate that then applies somewhere else, which starts a process of new applications

and rejections in the market. This can lead to a match with a partner who is preferred over the one the agent rejected. Due to uncertainty, this process can either lead to a more preferred candidate for the agent or fail and cause instability. Fernandez, Rudov, and Yariv (2022) show that the results of Roth (1989) hold with only minimal uncertainty on the proposing side of the market in a DA mechanism. In contrast, in our model, instability arises from learning about the preferences of others, which influences the expected utility of being matched with an agent. This implies that our findings are robust to the market side that faces uncertainty and hold, irrespective of strategic play.

Most closely related to our work is a recent literature of incomplete information and interdependent preferences, where match utilities depend on the type of the agent one is matched with. Hence, in contrast to Roth (1989), players' types affect their desirability and not just their preference reporting strategies. Both Chakraborty, Citanna, and Ostrovsky (2010) and Liu et al. (2014) build on the idea that agents can draw insights from the actions of other players. In Chakraborty, Citanna, and Ostrovsky (2010), each school receives a signal about the quality of the students on the other market side before submitting their preferences to the mechanism. By observing the final matching, a school can learn about the signals that other schools received and update their belief about the quality of a student. Anticipating that rematching is possible, the school is tempted to strategically misrepresent its signal, which can lead to instability. Liu et al. (2014) analyze a setting with transferable utility in which firms and workers are already matched. A firm only knows the quality of their matched worker. Firms can still update their information about the quality of other workers by observing rematching (or absence of rematching) in the market. In our model, we focus on the idea that types of agents differ in their preference profiles and that other agents care about these. Hence, agents may update their beliefs about the underlying types by only observing the final matching. Our stability results do not require anticipated rematching, nor do preferences of the own market side affect the desirability of other agents.

The rest of this paper is structured as follows. Section 3.2 outlines an illustrative example that provides intuitions for the formulation of reciprocal preferences and their consequences in matching markets. In Section 3.3, we set up the formal matching market and introduce reciprocal preferences. Section 3.4 analyzes the implications of reciprocal preferences in the DA mechanism, we then generalize our findings to a broader class of mechanisms in Section 3.5. In Section 3.6, we extend our analysis to a school choice setting with one non-strategic market side. Finally, Section 3.7 discusses the implications of our findings and concludes.

3.2 Illustrative Example

In this section, we provide intuitions for the consequences of reciprocal preferences on stability in matching markets through an illustrative example before presenting the main theoretical analysis. The example shows that the DA mechanism leads to an unstable allocation.

We consider a small job market with three firms (A, B, C), three workers (I, II, III), and a DA mechanism to match them one-to-one. All firms (A, B, C), as well as workers II and III have standard preferences. Both workers II and III prefer to only work for either of the firms over being unmatched. Worker II wants to work only for firm A , and worker III wants to be matched only with firm C . Worker I has reciprocal preferences: she⁶ cares how she is ranked by a firm. She prefers working for firm A over firm B if firm A ranks her first. If firm A ranks her second, then she prefers firm B over firm A . This means that her preference list is given by $A_1 \succ B \succ A_2$. The indices denote the true rank assigned to her by the respective firm. The (reciprocal) preferences of workers are common knowledge. Although the type of a firm is private knowledge, every agent knows the distribution of firms' types.

⁶We deviate from to convention to refer to all agents as "they/them" in the examples to make the examples more understandable and to avoid misunderstandings. We refer to institutions (firms/schools) as "they/them", to individuals (workers/students) as "she/her

Acquiring information about the true preferences of firms is challenging in practice. If the firms' preferences were perfectly observable, then the preference list of worker I would reduce to the standard case. Worker I prefers $A \succ B$ if firm A ranks her first, while she prefers $B \succ A$ if firm A ranks her second. In reality, potential employees typically only have limited information about the exact demands of a firm and face uncertainty about the characteristics of the competing applicants. Moreover, employers may not be interested in truthfully revealing their preferences so that they can give each applicant the impression that they are a preferred candidate. Hence, we allow for uncertainty about the firms' preferences.

We incorporate uncertainty about the firms' preferences as follows (we refer to the different realizations of preferences as *types*). A firm knows its own realized type, but the other agents do not. In this example, firm A has two possible types denoted by a superscript A^1, A^2 . Firm A^1 considers only worker I and II as potential employees, and has preferences of $I \succ II$. When being of type A^2 , it only considers workers III and I , and prefers $III \succ I$. The probability of firm A being of type A^1 is p . Firm B only wants to be matched with worker I and firm C only wants to be matched with worker III . We summarize the information on the matching market in Example 3.1. In addition, we assume that worker I has a higher expected utility of being matched with firm B than taking the lottery of being matched with firm A without knowing the type of firm A ($I : u(B) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$).⁷

Given their knowledge about the matching market and the mechanism, workers can infer the type of a firm after observing the final matching. For example, if firm A is matched with worker II , then agents can infer that firm A is of type A^1 because type A^2 does not consider worker II as a relevant candidate.

To illustrate the main intuitions, we first derive the optimal strategy of

⁷We only denote the preferences that a firm or worker has over being unmatched (e.g., the preferences $A : I \succ II \succ A \succ III$ are only denoted as $A : I \succ II$).

Table 3.1: Example 1

Proposer / Firm		Receiver / Worker
$A^1 : I \succ II$	with (p)	$I : A_1 \succ B \succ A_2$
$A^2 : III \succ I$	with $(1 - p)$	$II : A$
$B : I$		$III : C$
$C : III$		

Given:

$$I : u(B) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$$

worker I in the DA mechanism, and then show that the outcome is unstable. Except for worker I , all agents in the matching market have standard preferences and will state these truthfully.⁸ Given that all agents except worker I submit true preferences to the mechanism, both types of firm A , as well as firm B , will always make an offer to worker I during the DA mechanism (firm A^2 will always be rejected by worker III , and will therefore make an offer to worker I). Worker I states her preferences based on her expected utility. Assuming that her utility of being matched with firm B is higher than the lottery of being matched with types A^1 or A^2 , then she states $B \succ A$.

Given that worker I states $B \succ A$, the type of firm A will be revealed through the final matching. If type A^1 is realized, then worker I is matched with firm B , and firm A is matched with worker II . Through observing the match of worker II and firm A , worker I can infer that firm A is of type A^1 . Therefore, worker I wants to be matched with firm A . Given that firm A and worker I want to be matched to each other mutually, the matching is unstable. This happens because information about the type of a firm is revealed through the mechanism and the resulting final matching. We call this notion *immediate stability*.

⁸We will show later that it is a weakly dominant strategy for proposers to state true preferences in a DA mechanism, even if receivers have reciprocal preferences. Because workers II and III only consider working for one firm, it is also a weakly dominant strategy for them to state their true preferences.

3.3 Model

Overview. We consider a matching market with workers on one side of the market, firms on the other side, and a mechanism to match them one-to-one. Our matching setup differs from the standard model in Gale and Shapley (1962) in two aspects. First, the realization of firms' types and preferences is private information. Second, workers have reciprocal preferences, and therefore they prefer partners who also like to be matched with them. Agents may care about the fundamental preferences of others according to belief-based and preference-based motives. Being a preferred candidate can be a signal about the match-specific value that the firm may be better informed about (Avery & Levin, 2010; Lee & Niederle, 2015). Second, workers may enjoy interacting with a firm that likes them (Montoya & Horton, 2012; Montoya & Insko, 2008), even if preferences do not signal differences in relationship productivity. Hence, we allow for more general preference profiles than in a standard setting.

Set-up. Two disjoint finite sets of agents are one-to-one matched, following the marriage model of Gale and Shapley (1962). We consider firms $F = \{A, B, C, \dots\}$ and workers $W = \{I, II, III, \dots\}$. We sometimes denote an arbitrary firm with f and an arbitrary worker with w . Firms have strict and complete preferences over workers. These preferences are represented by an ordered list $P(f)$ on the set of $W \cup \{f\}$. A firm's preferences are private knowledge to the firm. Equivalently, we may think about the type (instead of the preferences) of firms being private information. The finite set of all possible types for a given firm is given by T_i . We denote the type of a firm by a superscript (e.g., $T_A = \{A^1, A^2, A^3, \dots\}$). The type of a firm f is independently drawn from a distribution g_f known to each agent in the market. Workers have reciprocal preferences over firms. The preference of a worker is represented by an ordered list $P(w)$ on the set of all possible types of all firms $T_A \cup T_B \cup T_C \cup \dots \cup \{w\}$. For example, a firm's preferences might be

$P(f) = II, I, f, III$. A worker's preference might be $P(w) = A^1, B^1, A^2, w, B^2$ denoting that a worker w prefers being matched with firm type A^1 over firm type B^1 over firm type A^2 over being unmatched f and over being matched with firm type B^2 . $\mathbf{P} = \{P(A), P(B), \dots, P(I), P(II)\dots\}$ denotes the set of all preferences. We also use the notation $i \succ j$ to state that i is preferred over j , both for firms and workers.

We examine the idea that a worker cares about being ranked well by the other market side. The true rank of worker w in the preference list of a firm is denoted by τ . Formally, we impose that an agent prefers $A_\tau \succeq A_{\tau+1}$. For example, an worker I weakly prefers to be the first choice of A than the second choice of A . By defining the match utility of worker I as $u_I(A_\tau)$, we generalize that if $i < j$ ($i, j \in \mathbb{N}$), then $u_I(A_i) \geq u_I(A_j)$. Being ranked better by firm A weakly increases the utility from being matched with firm A . This implies that we rule out cases in which agents like to be an undesirable alternative by a potential partner.⁹

In a matching market, each agent faces the decision about what preferences order $Q(f)$ (or $Q(w)$) to state given a mechanism h . $\mathbf{Q} = \{Q(A), Q(B), \dots, Q(I), Q(II)\dots\}$ is the set of all stated preferences by firms and workers. Following our theoretical set-up, a matching market is described by the agents in the matching market, their possible types and probabilities of realisation for every type, and their preference profile (including reciprocal preferences).

A matching mechanism h takes the stated preference profiles \mathbf{Q} and then maps them into a matching μ . A matching μ is a one-to-one correspondence. After the matching mechanism takes place, every firm f of the market is either matched with a worker denoted by $\mu(f) = w$ or is matched with itself $\mu(f) = f$. The same applies to workers.

⁹This may neglect potential behavioral mechanisms where a worse rank leads to higher desirability. For example, a worse rank may increase the desire to work with this party to convince it about one's quality. Alternatively, a better rank may be a signal of the worse quality of the other party [e.g., "*I don't care to belong to any club that will have me as a member.*" (Groucho Marx)]. Although theoretically possible, we consider these mechanisms to be secondary to a preference for a partner who likes to be matched with one.

Stability. In a setting of incomplete information, the stability criterion has to specify under which circumstances an agent would like to rematch. We use a standard framework of expected utility to define stability. An agent wants to rematch if their expected utility of rematching is higher than their expected utility of staying with their current match. Utility is non-transferable, and agents rematch if their (expected) utility from rematching is larger than their (expected) utility from their current matching. Defining stability in terms of expected utility corresponds to the concept of *Bayesian stability* in Bikhchandani (2017).¹⁰

Definition 3.1. *Bayesian stability:* A matching μ is Bayesian blocked by a worker-firm pair (A, I) that are not matched, but the expected utility of worker I and the utility of firm A increase by matching with each other. A matching μ is Bayesian blocked by an agent B (firm or worker) if the agent B prefers to be unmatched to their current match in (expected) utility. A mechanism is Bayesian stable if at least one of its equilibria is not blocked by any individual or any pair of agents for every realization of all possible type realizations T_i .

Our main stability concept evaluates *Bayesian stability* directly after the mechanism determines the matching, and the matching becomes public, which we call *immediate stability*. Once the matches are public, Bayesian updating about the types of other agents is possible. If this updating process leads to a blocking individual or blocking pair, then the matching is not *immediate stable*. Agents only infer information about types and preferences from the observed matching, which is in contrast to (for example) Roth (1989) where preferences become public knowledge.

Definition 3.2. *Immediate stability:* A matching is immediate stable if it is Bayesian stable after the outcome of the matching mechanism is known to the

¹⁰Other papers consider different stability concepts; for example, the idea that a matching is not stable if there is a positive probability to profit from rematching (Lazarova & Dimitrov, 2017), when there is no probability that the rematching leads to a worse outcome (ex-ante stability in Bikhchandani, 2017), or whenever there is scope for a mutually beneficial outcome given that transfer payments are possible (Liu et al., 2014).

agents.

This definition of stability allows us to characterize matching mechanisms that lead to stable outcomes. A mechanism guarantees stability if at least one of its equilibria is always Bayesian stable. This implies that the mechanism leads to stable outcomes for every realization of the firms' types. Hence, a mechanism is considered to be stable when every possible outcome of the given equilibrium is stable. It is important to note that we implicitly define stability together with the underlying mechanism because of potentially different degrees of information revelation. While specific outcomes may be stable with one mechanism, they are not stable with another.¹¹

For our secondary stability concept, we assume that a player learns the type of their matched agent after some time. We consider it reasonable that the type of a partner gets revealed through interaction. Therefore, this constitutes another natural point to evaluate stability. If an agent wants to rematch after seeing the final matching and learning the true type of their partner, then the matching is not *ex-post stable*. Hence, *immediate stability* and *ex-post stability* are evaluated at different points in time (see Figure 3.1) but rely on the same stability concept (*Bayesian stability*). At both points in time, agents update their beliefs about the firms' types.

Definition 3.3. *Ex-post stability: A matching is ex-post stable if it is Bayesian stable after the outcome of the matching mechanism is known to the agents and workers learn the true type of the firm that they are matched with.*

There is a connection between both stability notions in that ex-post stability implies immediate stability. The proof of Proposition 3.1 is delegated to the Appendix (see Section C.1). The intuition is that an immediate outcome can result in one or more ex-post outcomes. If all of the ex-post outcomes are

¹¹For illustration, let us assume that a mechanism randomly pairs workers and firms without taking the stated preferences into account. This mechanism may lead to the same outcome as a DA mechanism. In the former cases, workers cannot infer something from the outcome about firms' true preferences, while they can in a DA mechanism.




Figure 3.1: Timeline

stable, then the immediate outcome must also be stable because a mechanism is stable if at least one of its equilibria is always stable with respect to the true preferences. Meanwhile, immediate stability does not guarantee ex-post stability.¹² Imagine that there is one ex-post unstable outcome that only materializes with a small probability. In this situation, it can still be the case that the outcome is immediate stable. Therefore, our main results focus on the more restrictive concept of *immediate instability*, any results related to *ex-post stability* are deferred to the Appendix.

Proposition 3.1. *In a matching market with reciprocal preferences, ex-post stability is a sufficient condition for immediate stability.*

Strategy-proofness. The concept of strategy-proofness generally carries over from the complete information setting. A mechanism is strategy-proof if it is optimal to report preferences truthfully for every agent and for all strategies of other agents. In our setting, strategy-proofness for a standard mechanism such as the DA mechanism is only defined for agents without reciprocal preferences because workers with reciprocal preferences cannot state their full preferences profile.

Definition 3.4. *Strategy-proofness: A mechanism is strategy-proof if truthful preference revelation (or simply truth-telling) is a weakly dominant strategy for every agent.*

¹²The same correspondence holds between the concept of ex-post stability and the complete information stability concept of Roth (1989), where all agents' types become public knowledge. In a matching market with reciprocal preferences, stability according to Roth (1989) is a sufficient condition for ex-post stability.

3.4 Two-Sided Matching Markets: Deferred Acceptance Mechanism

The DA mechanism (Gale & Shapley, 1962) plays a key role in two-sided markets because it achieves stable matching outcomes in a standard complete information setting. Agents on both sides of the market (proposers and receivers) submit a rank-ordered preference list about the acceptable agents on the other side of the market to the mechanism. In the DA mechanism, stability is compatible with truth-telling for the proposing side of the mechanism. Under complete information, the equilibrium misrepresentation of preferences by the receiving side still results in stable outcomes (Roth, 1984). Given that there is no stable and strategy-proof mechanism for both sides of the market (Roth, 1982), the DA mechanism comes close to achieving the optimal outcome.

Due to these desirable properties, we first analyze reciprocal preferences in the standard DA mechanism before generalizing our main findings to a broader class of mechanisms. In the DA mechanism, both sides of the market must submit a standard preference list to the mechanism in which each acceptable agent of the other market side appears only once. Therefore, an agent with reciprocal preferences must submit a standard preference list. For example, a worker with reciprocal preferences $A_1 \succ B \succ A_2$ has to decide whether to state $A \succ B$ or $B \succ A$ (assuming that being matched is better than remaining unmatched).

In our setting, truth-telling remains a weakly dominant strategy for proposers without reciprocal preferences. This follows immediately from the results in the standard setting of the DA mechanism. While the settings with reciprocal preferences may change the strategy of other agents, the mechanism remains the same. By the definition of a weakly dominant strategy, the potential adjustments of other players will not change the optimality of stating true preferences for the proposer.

Proposition 3.2. *Truth-telling is a weakly dominant strategy for proposers with*

standard preferences in a one-to-one-matching under the DA mechanism.

Proof. Suppose not. Then there must be another strategy of the proposer which does better for at least one set of played strategies by the other players. But this contradicts the fact that, given any set of strategies by the other players, truth-telling is a weakly dominant strategy in a DA mechanism. \square

Meanwhile, receivers may have an incentive to misrepresent their preferences. This directly follows from the fact that our framework encompasses the standard case, in which truth-telling is not a weakly dominant strategy for receivers. Following standard assumptions, we do not consider a rematching stage in the market. In Appendix C.2.1, we relax this assumption and assume that agents anticipate the rematching stage. Rematching leads to strategic misrepresentation of preferences by proposers in the first stage.

In the next step, we analyze the effect of reciprocal preferences on stability. As shown through Example 3.1, the DA mechanism ceases to be immediate stable if agents have reciprocal preferences. This is even true if receivers do not misreport their preferences. We additionally show that there are matching markets where no strategy for a worker with reciprocal preferences leads to an immediate stable matching. This means that a receiver with reciprocal preferences cannot choose any strategy that guarantees immediate stability, even though all of the other agents state true preferences. In the case of standard preferences, an outcome is always stable if every agent states their true preferences (Gale & Shapley, 1962). A similar result cannot be established in the case with reciprocal preferences.

We demonstrate that worker I cannot choose any strategy that guarantees an immediate stable outcome by analyzing the same market as in Example 3.1. Depending on the stated preferences of worker I , worker I can be matched with firm A , firm B or remain unmatched. Every preference submission by worker I has to result in one of these outcomes. We show that each of these three outcomes can be immediate unstable. We have already shown that being

matched with firm B (e.g., by stating $I : B$ or $I : B \succ A$) is immediate unstable because worker I infers the type of firm A after seeing the final matching. Worker I wants to be matched with firm A after learning that the firm is of type A^1 . If worker I is matched with firm A (e.g., by stating $A \succ B (\succ I)$ or by just stating $A (\succ I)$), then the matching will always be immediate unstable. Worker I does not learn the type of A and forms a blocking pair with firm B because she prefers being matched with firm B over firm A if she does not know the type of firm A ($I : u(B) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$). In the last case where worker I remains unmatched (by not stating a firm), she forms a blocking pair with either of the two firms. This shows that worker I cannot submit any preference profile that guarantees stability due to the information revelation by seeing the final matching.

Proposition 3.3. *The DA mechanism with reciprocal preferences is not always immediate stable. There are markets where there is no strategy for a worker that leads to an immediate stable outcome if all the other players state true preferences.*

In Appendix C.2.2, we show that there are matching markets that are always ex-post unstable for every possible strategy that a player can choose in the DA mechanism. At the same time, if agents were free to rematch directly after the matching, then there are ex-post stable outcomes that agents would not reach because of immediate instability. A worker may rationally want to leave her match after the matching (immediate), but—given the realized type of her match—still prefers to be matched with this firm (ex-post). We show this in Appendix C.2.3. This implies that, even though ex-post stability is a sufficient condition for immediate stability, transparency about the outcomes of the matching mechanism prevents certain stable matching outcomes from materializing if agents are free to rematch.

Instability also arises when the agents with reciprocal preferences are on the proposing side and the agents with types are on the receiving side of the

DA mechanism. Our main findings of instability due to reciprocal preferences apply irrespective of the market side that faces uncertainty. This implies that information about which side of the market has reciprocal preferences cannot solve the problem from the designer's perspective. The following Example 3.2 shows that the DA mechanism is neither immediate nor ex-post stable with reciprocal preferences on the proposing and uncertainty on the receiving side.

Table 3.2: Example 2

Proposer / Worker	Receiver / Firm
$I : A_1 \succ B \succ A_2$	$A^1 : I$ with (p)
$II : II$	$A^2 : II \succ I \succ III$ with $(1 - p)$
$III : A$	$B : I$

Given:

$$I : u(B) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$$

Worker I has to decide about stating $A \succ B$ or $B \succ A$. If she states $A \succ B$, then she will be matched with firm A ; and if she states $B \succ A$, then she will be matched with firm B .¹³ We assume that worker I prefers to be matched with firm B rather than being matched with firm A without knowing the type [$u(B) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$]. Hence, she states $B \succ A$ and is matched with firm B . However, worker I will infer the type of firm A through seeing the final matching. If firm A is of type A^2 , then it is matched with worker III , while firm A remains unmatched if it is of type A^1 . In the case of firm A^1 , worker I and firm A want to rematch. The matching is immediate and ex-post unstable.

Regarding agents' incentives, we show that proposers do not have a weakly dominant strategy in this setting. While in a standard setting of uncertainty, truthful preference revelation is a weakly dominant strategy for the proposing side, there is not one in the case of reciprocal preferences. The optimal

¹³All other players state true preferences in equilibrium.

strategy of a proposer with reciprocal preferences depends on the behavior of other market participants. We show this in a simple example (Appendix C.2.4). Therefore, strategic considerations may even play a role for the proposing side in an environment where proposers have reciprocal preferences.

3.5 Stability in Two-Sided Matching Markets

This section analyzes whether an alternative mechanism can remedy the observed instability under the DA mechanism. We generalize the findings and obtain an impossibility result: there is no mechanism that is immediate or ex-post stable. We first demonstrate that no mechanism always achieves immediate stability. For this, it is sufficient to show that a matching market exists for which no mechanism can achieve immediate stability.

Proposition 3.4. *There is no mechanism that is always immediate stable for every matching market with reciprocal preferences.*

We demonstrate this by showing one matching market where no mechanism can achieve immediate stability.

Proof. Let h be a matching mechanism. The matching mechanism h selects a matching μ for the stated preference profiles \mathbf{Q} . A matching μ is a one-to-one correspondence and denotes with whom agents are matched. We show that no mechanism h can select an immediate stable matching according to the true preferences in the following matching market.

The market consists of firm A and three workers (I, II, III). Firm A can have two different types (A^1, A^2). Worker I has reciprocal preferences, while workers II and III have standard preferences. The preference profiles \mathbf{P} are defined in Example 3.3. Worker I prefers to be unmatched rather than being matched with firm A given their prior about firm A 's type in this example ($I : u(I) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$).

Table 3.3: Example 3

Firm		Worker
$A^1 : I \succ II \succ A \succ III$	with (p)	$I : A_1 \succ I \succ A_2$
$A^2 : III \succ I \succ A \succ II$	with $(1 - p)$	$II : A \succ II$
		$III : III$

Given:

$$I : u(I) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$$

The proof shows that no mechanism can select a match for both types A^1 and A^2 of firm A that is always immediate stable. Every possible matching is either immediate unstable because the firm's type is revealed and a blocking pair emerges or because one of the firm types has an incentive to mimic the other, which leads to immediate instability.

We proceed in three steps to establish that no mechanism is immediately stable. First, we exclude from the set of potentially stable matchings those that are immediate unstable without any information update and given worker I 's prior about the firm's type. Second, we narrow down the set of immediate stable matchings by eliminating all matchings where the information updating about the firm's type through seeing the final matching leads to instability. Third, we show that the remaining set of possibly stable matching cannot be reached due to strategic play of the firm and a resulting pooling equilibrium that is always immediate unstable.

Step 1: To reduce the set of potentially stable matchings, we first exclude all matchings that are unstable given the preferences of agents and the workers' priors about firm A 's type. A matching can only be immediate stable if firm A^1 is matched with worker I or II , and type A^2 is matched with worker I or remains unmatched. A matching is immediate unstable in all the other cases. Worker III prefers to be unmatched regardless of firm A 's type, while firm A prefers to be unmatched if it is of type A^2 and gets matched to worker II and

firm A^1 prefers to be matched with worker II if being unmatched.

Step 2: We further narrow down the set of immediate stable matchings by considering the information updating of workers through seeing the final matching. We start on the basis of the remaining possible, stable matchings after Step 1 and show that firm A^1 cannot be matched with worker II . If there is a positive probability that both type's of firm A are matched to worker II , then the matching is immediate unstable because firm A^2 prefers being unmatched over being matched with firm A^2 . Hence, if the mechanism is immediate stable it can only match firm A^1 with worker II . However, if firm A^1 is matched with worker II , worker I can infer the type of firm A^1 . The matching is immediate unstable because firm A^1 and worker I form a blocking pair. Therefore, given a immediate stable mechanism, firm A^1 cannot be matched with worker II and can only be matched with worker I .

Step 3: We show that every remaining possible, stable matching cannot be reached due to the strategic play of firm type A^2 . For every matching that has a positive probability of firm A^2 being unmatched and firm A^1 being matched with worker I for sure, type A^2 has an incentive to mimic type A^1 . This would result in the matching $\mu(A^1) = I$ and $\mu(A^2) = I$, which is not immediate stable by assumption because worker I prefers to be unmatched compared to being matched with firm A without knowing the type of firm A ($I : u(I) > p \cdot u(A_1) + (1-p) \cdot u(A_2)$). This implies that every possible matching is immediate unstable. \square

To demonstrate that there is no mechanism that always leads to an ex-post stable matching, we build on Proposition 3.1.

Proposition 3.5. *In a setting with reciprocal preferences, no mechanism always leads to an ex-post stable matching.*

Proof. Proposition 3.1 states that ex-post stability is a sufficient condition for immediate stability in a matching market with reciprocal preferences. We show in Proposition 3.4 that there is no immediate stable mechanism. By law of

contraposition, a matching cannot be ex-post stable if it is immediate unstable. Therefore, we can conclude that there is no ex-post stable mechanism. \square

3.6 School Choice Markets

In the school choice setting, only one side of the market consists of strategic agents (Abdulkadiroğlu & Sönmez, 2003). The other side has priorities that are not subject to misrepresentations. This framework especially applies to public institutions (e.g., schools), which have priorities over other agents (e.g., students). These priorities must be truthfully reported to the mechanism. For example, a school may be required by law to prioritize its applicants according to specific rules (e.g., scores in entrance exams or distance to the students' residence). Understanding the effects of reciprocal preferences in school choice is crucial in its own right but it can also provide insights for effective mechanism design when comparing outcomes with those of standard two-sided markets. We will show that excluding strategic considerations for one side of the market, combined with choosing the right mechanism, resolves the instability associated with reciprocal preferences.

Like the standard two-sided market, reciprocal preferences only affect stability in a school choice setting when there is uncertainty about the priorities. One may think about (at least) three different reasons for uncertainty in the school choice setting. First, the rules on how priorities are formed might not be publicly available. Second, the rules are publicly available but the costs to understand them are high, which may be especially true when considering a sizeable potential choice set. Third, an applicant knows how priorities are formed but lacks information about the characteristics of other applicants, and she therefore faces uncertainty about the final priorities.

Under uncertainty about the priorities, the DA mechanism results in unstable outcomes. The uncertainty about schools' priorities can still result in ex-post suboptimal decisions of the applicants. The intuition is that applicants

have to make a choice under uncertainty about the true priorities because these are not revealed ex-ante. Given that individuals care about their true ranking in the preference lists, this can result in unstable matchings. We show this through Example 3.4, where applicant I has to decide between stating $A \succ B$ or $B \succ A$. Given that uncertainty applicant I prefers being matched with school B rather than being matched with school A given her prior about the schools' type [$I : u(B) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$], she states $B \succ A$. Applicant I learns whether school A is of type A_1 or A_2 after seeing the final matching. If firm A is unmatched worker I knows that it is of type A^1 and they build a blocking pair. When school A is of type A^1 , the outcome is immediate and ex-post unstable.¹⁴

Table 3.4: Example 4

Proposer / School		Receiver / Applicant
$A^1 : I$	with (p)	$I : A_1 \succ B \succ A_2$
$A^2 : II \succ I \succ III$	with ($1 - p$)	$II : II$
$B : I$		$III : A$

Given:

$$I : u(B) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$$

We present two simple remedies to overcome this inefficiency: first by inducing information revelation before applicants submit their preference list, and second by incorporating reciprocal preferences into the ranking. We consider a simple sequential variant of the DA mechanism to induce information revelation. In this Two-Stage Deferred Acceptance (TSDA) mechanism, the market side with uncertain preferences first submits their rankings publicly to the mechanism. The side with reciprocal preferences then submits their preference ranking, incorporating the information from the first step. With these preferences, we run a standard DA mechanism. Therefore, the only difference between

¹⁴In Appendix C.2.5, we show that this does not depend on whether the applicants are on the receiving or proposing side of the algorithm.

a standard DA and the TSDA mechanism is the preference submission's timing and observability.

The TSDA mechanism solves the information asymmetries without any caveats in the school choice problem. Institutions state their priorities truthfully, which results in a situation of complete information. Reciprocal preferences can affect the preference order but every applicant can rank schools unambiguously by incorporating the information about the schools' preferences. Truth-telling then means submitting those preferences. If applicants are on the receiving side of the algorithm, they might still misrepresent their preferences in equilibrium as in the standard DA mechanism but the outcome is stable (see Roth, 1984). If applicants are on the proposing side, then they do not have any incentive to misrepresent their preferences. Hence, a TSDA mechanism with the applicants as the proposers leads to truth-telling and stability.

A variant of the standard DA mechanism that allows individuals to submit their complete preferences order (even if they have reciprocal preferences) also leads to stability. In the standard DA mechanism, participants submit a rank-ordered list in which all of the acceptable partners are listed once. Instead, the preferences lists of our modified version of the DA mechanism can include reciprocal preferences. This means that agents can submit their full contingent preference lists to the mechanism (e.g., $I : A_1 \succ B \succ A_2$ instead of having to decide whether to rank A over B , or vice versa). We call this version the Reciprocal References Deferred Acceptance (RPDA) mechanism.

The mechanism then uses the (simultaneously) stated preferences of all agents and creates a standard preference ranking for every agent. For applicant I who submitted the preference profile $A_1 \succ B \succ A_2$ to the mechanism, the mechanism assesses whether applicant I prefers school A over B based on the stated preferences of these schools. If school A stated applicant I as their highest preference, then the mechanism assigns the preference profile $A \succ B$ to agent I . If school A did not state applicant I as their highest preference, then the mechanism uses the ranking $B \succ A$. With these rank-ordered lists,

a standard DA mechanism takes place. Implicitly, the mechanism takes the stated preferences as the true preferences. While the reciprocal preferences of agents are based on the true preferences, the mechanism determines the final rankings by using the stated preference information.

The RPDA mechanism is stable when applicants are on the proposing side of the algorithm (and schools with standard preferences on the receiving side). The proposing side has a weakly dominant strategy to state true preferences and the receiving side submits their (true) priorities.¹⁵ Applicants do not have to take into account any strategic considerations when stating their complete preference profile. Like the TSDA mechanism, the RPDA mechanism is strategy-proof and stable when individuals are on the proposing side. However, when applicants are on the receiving side of the mechanism, uncertainty and preference misrepresentations can lead to instability. As in standard DA mechanism, the RPDA mechanism is not strategy-proof for the receiving side. By misrepresenting their preferences, receivers might be able to implement the receiver optimal stable matching (instead of the proposer optimal matching). Due to the uncertainty, these misrepresentation can result in immediate and ex-post unstable outcomes, even in the absence of reciprocal preferences.

Proposition 3.6. *The DA mechanism does not achieve stability in a school choice setting where applicants have reciprocal preferences. In contrast, a sequential variant of the DA (TSDA) mechanism and a variant in which agents can submit their reciprocal preferences to the mechanism (RPDA) can achieve stability.*

Therefore, the TSDA or the RPDA mechanisms may be perceived as solu-

¹⁵Stating true preferences is a weakly dominant strategy for proposers in the standard DA mechanism (Roth, 1982). Therefore, an agent cannot do better than stating (reciprocal) preferences in the RPDA mechanism that correspond to these standard preferences. In the school choice setting, schools are non-strategic and reveal their priorities truthfully. Given that applicants state their true reciprocal preferences, the standard preference ranking produced by the RPDA mechanism precisely reflects the applicants' preferences under complete information. Therefore, stating true reciprocal preferences is a weakly dominant strategy in the RPDA mechanism.

tions to situations with reciprocal preferences in standard two-sided matching markets with strategic players on both sides of the market. We demonstrate in Section 3.5 that there are no mechanisms that achieve *Bayesian stability* (both after agents observe the mechanism's outcome and once agents learn the type of their partner). This implies that neither the TSDA nor the RPDA mechanism can guarantee *Bayesian stability*.

In contrast to the DA mechanism, neither TSDA and RPDA mechanisms are strategy-proof if we consider a standard setting with strategic firms on the proposing side (see the proof in Appendix C.2.6). The underlying idea of the TSDA and RPDA mechanism is that workers can react to the types and preferences of firms. This also implies that firms start to send favorable signals to applicants strategically. For example, despite preferring a very talented applicant who is extremely unlikely to join a (mediocre) firm in any case, this firm may use its top signal for an applicant whose decision could be influenced by it. One can even show that there are markets in which the DA mechanism achieves stable outcomes in undominated strategies that imply truthful reporting for firms, while the TSDA mechanism does not (see Appendix C.2.7).

3.7 Discussion and Conclusion

In this paper, we motivate, formalize, and analyze the effects of reciprocal preferences in matching markets. Reciprocal preferences allow for the possibility that preferences of the other market side influence an agent's preferences. We provide three main results when agents care about others' preferences without perfectly knowing these. First, we show that the DA mechanism ceases to be stable when agents observe the final allocation of the mechanism. Even if agents do not strategically misrepresent their preferences, the final matching can be unstable. Second, we demonstrate that no mechanism always leads to a stable matching when (at least) one of the agents has reciprocal preferences. Third, when extending our analysis to a school choice setting, we show that variants

of the DA mechanism achieve stability.

We assume one-sided reciprocal preferences throughout this paper. In principle, this can be extended to consider reciprocal preferences on both sides of the market. Once we do so, we need to define how the reciprocal preferences of worker I correspond to firm A 's preferences if these are also reciprocal to avoid a recursive problem. If firm A 's preference order is $I_1 \succ II_1 \succ I_2 \succ II_2$, then it is unclear whether worker I is ranked better or worse than worker II in the preference list of firm A (and hence how worker I 's preferences should respond to these). One reasonable possibility to solve this problem is to assume that every reciprocal preference order must have a general structure, such that being ranked the same by every agent on the other market side induces the same ranking. Given the reciprocal preferences of firm A , this *general preference order* would be $I \succ II$, because the firm prefers $I \succ II$ when being ranked first by both workers, and when being ranked second by each worker. Once we assume that there exists such a *general preference order* and that reciprocal preferences are based on it, the model becomes tractable again.

Understanding the effects of reciprocal preferences is critical for understanding why certain matching markets do not perform satisfactorily. First, reciprocal preferences provide a rationale for why agents may prefer decentralized matching over centralized (algorithmic) matching mechanisms. Decentralized markets in which agents interact allow them to learn the other sides' preferences, which makes it hard to establish a centralized mechanism if they care about others' preferences. Second, modifications of the mechanism by participants can be a consequence of reciprocal preferences. Agents may add screening tools to attract especially interested candidates, such as early admission (Avery & Levin, 2010), or may even base their own ranking directly on other agents' preferences (Chiu & Weng, 2009; U. Dur et al., 2022). Third, the results help us to understand observed instability in centralized matching markets. Given that true preferences are largely unobservable, theory can help evaluate the reasons for instability and allow for a better design of matching mechanisms. Policy con-

clusions given observed instability in a mechanism crucially depend on whether the reason for instability is a strategic misrepresentation of preferences or are reciprocal preferences per se.

This paper informs market designers about what information should be disclosed. We show that information about others' preferences is crucial for stability. While the market designer has arguably little influence on the agents' learning preferences of others in individual conversations, they can guide information flows by revealing all final matches to the agents (or even the submitted preference rankings). This is not only crucial for the effect of reciprocal preferences to come into play but it also matters if agents have other non-standard preferences, such as regret that crucially depends on counterfactual outcomes.

Our findings on stability in school choice also inform policy where institutions are strategic agents and forcing them to act non-strategic might be impossible. In this case, one may oblige schools to determine preferences based on objective and transparent criteria. The need for schools to determine their ranking based on these criteria increases the credibility that the ranking corresponds to the actual preferences. Ranking applicants according to some pre-specified criteria where compliance can be (partially) verified by the market designer mitigates the scope of submitting non-truthful preferences and enhances stability. In that sense, designing rules for a matching market that make it more similar to a school choice setting can increase stability and welfare.

This paper highlights that standard matching mechanisms do not function as desired when agents care about others' preferences without perfectly knowing them. This is true not only for reciprocal preferences but also for more generally situations in which agents have type-dependent preferences, with the types themselves being defined by the preferences of the other agents. While this paper focuses on reciprocal preferences, investigating different classes of type-dependent preferences in matching markets remains for future research. Natural extensions of our current analysis include considering additional information structures about initial preferences, other information sets when eval-

uating stability, and different stability concepts under uncertainty. While we derive our (negative) results on stability through the theoretical analysis of markets, we have a limited understanding of how often instability actually occurs in markets when using standard mechanisms. Therefore, in complementary work (Opitz & Schwaiger, 2022a), we show the relevance of reciprocal preferences for (in)stability through a laboratory experiment and validate our underlying theoretical assumption.

Appendix A

Appendix to Chapter 1

A.1 Proof of Proposition 1.2

Proof. This proof solves for a symmetric Bayesian Nash Equilibrium bidding strategy $b_i(v) = b(v) \quad \forall i = 1, \dots, n$ in a FPSB auction. A manager wins the auction if he submits the highest bid and has to pay a price equal to his bid $b_i(v)$. The expected payoff of a manager i is: $U_i(b_i, b_{-i}, v_i, a, s) = Pr[b_j \leq b_i, \forall j \neq i] \cdot (s(v_i - b_i) + a)$.

Bidder i chooses $b_i(v)$ to solve:

$$\max_{b_i} u : \left(F(b^{-1}(b_i)) \right)^{n-1} \cdot (s(v_i - b_i) + a) \quad (\text{A.1})$$

F.O.C.:

$$(s(v_i - b_i) + a)(n - 1) \left(F(b^{-1}(b_i)) \right)^{n-2} f(b^{-1}(b_i)) \frac{1}{b'(b^{-1}(b_i))} - s \left(F(b^{-1}(b_i)) \right)^{n-1} = 0$$

At a symmetric equilibrium $b_i = b(v_i)$. The F.O.C reduce to:

$$(s(v - b(v)) + a) \cdot (n - 1)F(v)^{n-2}f(v)\frac{1}{b'(v)} = sF(v)^{n-1}$$

$$(s(v - b(v)) + a) \cdot (n - 1)\frac{f(v)}{F(v)} = b'(v)s$$

The next step solves the problem for the uniform distribution on $[0, 1]$: $F(v) = v$ and $f(v) = 1$

$$b'(v) \cdot s = (s(v - b(v)) + a) \cdot (n - 1)\frac{1}{v} \quad (\text{A.2})$$

$$b'(v) \cdot v = (n - 1)(v + \frac{a}{s}) - (n - 1) \cdot b(v)$$

$$b'(v) \cdot v + (n - 1) \cdot b(v) = (n - 1)(v + \frac{a}{s})$$

$$b'(v) \cdot v^{n-1} + (n - 1) \cdot b(v)v^{n-2} = (n - 1)(v^{n-1} + \frac{a}{s}v^{n-2})$$

$$\frac{\partial}{\partial v} (b(v) \cdot v^{n-1}) = (n - 1)(v^{n-1} + \frac{a}{s}v^{n-2})$$

$$\int_0^v \frac{\partial}{\partial t} (b(t) \cdot t^{n-1}) dt = (n - 1) \int_0^v t^{n-1} + \frac{a}{s}t^{n-2} dt$$

$$b(v) \cdot v^{n-1} = (n - 1) \left(\left[\frac{1}{n}t^n \right]_0^v + \frac{a}{s} \left[\frac{1}{n-1}t^{n-1} \right]_0^v \right)$$

$$b(v) \cdot v^{n-1} = \frac{n-1}{n} \cdot v^n + \frac{a}{s} \left(\frac{n-1}{n-1} \right) v^{n-1}$$

$$b(v) = \frac{n-1}{n} \cdot v + \frac{a}{s}$$

The bidding function is equal to the standard bidding function $b_{Standard}^{FPSB}(v) = \frac{n-1}{n}v$ plus the fraction $\frac{a}{s}$.

Next, we show that this is an equilibrium that maximizes utility. We assume a bidder can bid like a bidder with type x while every other bidder follows the derived bidding strategy. We show that it is optimal for a bidder not to deviate ($x = v$).

The expected utility of a bidder is $u(x, v) = x^{n-1}(a + s(v - b(x)))$

$$\begin{aligned} \max_b u &: x^{n-1} \left(a + s \left(v - \frac{n-1}{n}x - \frac{a}{s} \right) \right) \\ &= x^{n-1}sv - x^n s \frac{n-1}{n} \end{aligned}$$

We take the derivative to solve for the optimal x .

$$\begin{aligned} \frac{\partial u}{\partial x} &= (n-1)x^{n-2}sv - nx^{n-1}s \frac{n-1}{n} = 0 \\ (n-1)x^{n-2}sv &= nx^{n-1}s \frac{n-1}{n} \\ v &= x \end{aligned}$$

We take the second derivative to show that this is indeed a maximum.

$$\frac{\partial^2 u}{\partial x^2} = (n-2)(n-1)svx^{n-3} - (n-1)^2sx^{n-2}$$

This is a maximum because the function is concave around $v = x$.

$$(n-2)(n-1)svx^{n-3} - (n-1)^2sx^{n-2} < 0$$

$$n-2 < n-1$$

□

A.2 Proof of Proposition 1.4

Proof. As shown, a manager will never bid higher than v if the potential punishment is harsh enough $z > a$. Managers are restricted to not bid higher than v . First, we show that up to value $\bar{v} = (n-1)\frac{a}{s}$ a manager bids according to $b = v$.

Given that all other bidders bid $b_j = v_j$ where $v_j < (n-1)\frac{a}{s} = \bar{v} \forall j$, if $v_i \leq \bar{v}$,

we want to show that a manager i also bids $b(v_i) = v_i$.

(1) If the bid $b > v$, the manager's utility is negative due to the punishment.

(2) Now we show that underbidding is not optimal when $v < (n-1)\frac{a}{s}$.

Suppose the manager bids $b = v - \epsilon > 0$. Then his problem is:

$$\begin{aligned} \max_{\epsilon} \quad & (v - \epsilon)^{n-1}[a + s\epsilon] \\ \text{s.t.} \quad & \epsilon \geq 0 \end{aligned}$$

F.O.C.:

$$(n-1)(v - \epsilon)^{n-2} \cdot (-1)(a + s\epsilon) + (v - \epsilon)^{n-1}s \leq 0$$

with " $=$ " if $\epsilon > 0$.

Suppose now, $\epsilon > 0$. Then,

$$\begin{aligned} (v - \epsilon)^{n-2}[-(n-1)(a + s\epsilon) + (v - \epsilon)s] &= 0 \\ \Rightarrow \quad \epsilon &= \frac{1}{n}\left[v - (n-1)\frac{a}{s}\right] \end{aligned}$$

But this is a contradiction because with $v \leq (n-1)\frac{a}{s}$ this results in $\epsilon < 0$. Manager i wants to deviate by bidding higher. However, this option is already ruled out. Therefore, $\epsilon = 0$.

If $\epsilon = 0$:

$$(n-1)(v)^{n-2}(-1)a + (v)^{n-1}s < 0$$

$$v^{n-1}s < (n-1)v^{n-2}a$$

$$vs < (n-1)a$$

$$v < (n-1)\frac{a}{s}$$

All bidders with a value lower than $v < \bar{v}$ bid according to the symmetric bidding function $b(v) = v$. Bidders with a value $v \geq \bar{v}$ follow a symmetric Bayesian Nash Equilibrium bidding strategy $b(v) = b(v)$. The maximization problem is the same as in the case without punishment (equation (A.1)). However the support is $\left[(n-1)\frac{a}{s}, 1\right]$ and it assumes that a bidder with value \bar{v} is bidding $b(\bar{v}) = \bar{v}$.

As these proofs follow the same first steps as those above, the following proofs begin with equation (A.2). Again, solving for the uniform distribution: $F(v) = v$ and $f(v) = 1$.

$$b'(v) \cdot s = (s(v - b(v)) + a) \cdot (n-1)\frac{1}{v}$$

$$b'(v) \cdot v = (n-1)\left(v + \frac{a}{s}\right) - (n-1) \cdot b(v)$$

$$b'(v) \cdot v + (n-1) \cdot b(v) = (n-1)\left(v + \frac{a}{s}\right)$$

$$b'(v) \cdot v^{n-1} + (n-1) \cdot b(v)v^{n-2} = (n-1)\left(v^{n-1} + \frac{a}{s}v^{n-2}\right)$$

$$\frac{\partial}{\partial v} (b(v) \cdot v^{n-1}) = (n-1)\left(v^{n-1} + \frac{a}{s}v^{n-2}\right)$$

$$\int_{\bar{v}}^v \frac{\partial}{\partial t} (b(t) \cdot t^{n-1} dt) = (n-1) \int_{\bar{v}}^v t^{n-1} + \frac{a}{s}t^{n-2} dt$$

$$b(v)v^{n-1} - b(\bar{v})\bar{v}^{n-1} = (n-1) \left(\left[\frac{1}{n}t^n \right]_{\bar{v}}^v + \frac{a}{s} \left[\frac{1}{n-1}t^{n-1} \right]_{\bar{v}}^v \right)$$

As shown before, the support is $\left[(n-1)\frac{a}{s}, 1\right]$ and therefore $\bar{v} = (n-1)\frac{a}{s}$ and a

bidder with valuation \bar{v} bids $b(\bar{v}) = b\left((n-1)\frac{a}{s}\right) = (n-1)\frac{a}{s}$.

$$\begin{aligned}
b(v)v^{n-1} - b\left((n-1)\frac{a}{s}\right)\bar{v}^{n-1} &= (n-1) \left(\left[\frac{1}{n} t^n \right]_{\bar{v}}^v + \frac{a}{s} \left[\frac{1}{n-1} t^{n-1} \right]_{\bar{v}}^v \right) \\
b(v)v^{n-1} - (n-1)\frac{a}{s} \cdot \bar{v}^{n-1} &= \frac{n-1}{n}v^n - \frac{n-1}{n}\bar{v}^n + \frac{a}{s}v^{n-1} - \frac{a}{s}\bar{v}^{n-1} \\
b(v) &= \frac{n-1}{n}v - \frac{n-1}{n}\frac{\bar{v}^n}{v^{n-1}} + \frac{a}{s} - \frac{a\bar{v}^{n-1}}{sv^{n-1}} + (n-1)\frac{a\bar{v}^{n-1}}{sv^{n-1}} \\
b(v) &= \frac{n-1}{n}v + \frac{a}{s} - \frac{n-1}{n}\frac{\bar{v}^n}{v^{n-1}} + \frac{a}{s}(n-2)\frac{\bar{v}^{n-1}}{v^{n-1}} \\
b(v) &= \frac{n-1}{n}v + \frac{a}{s} - \frac{n-1}{n}\frac{\left((n-1)\frac{a}{s}\right)^n}{v^{n-1}} + \frac{a}{s}(n-2)\frac{\left((n-1)\frac{a}{s}\right)^{n-1}}{v^{n-1}} \\
b(v) &= \frac{n-1}{n}v + \frac{a}{s} - \frac{(n-1)^{n+1}\left(\frac{a}{s}\right)^n}{nv^{n-1}} + (n-2)\frac{(n-1)^{n-1}\left(\frac{a}{s}\right)^n}{v^{n-1}} \\
b(v) &= \frac{n-1}{n}v + \frac{a}{s} - \frac{\left(\frac{a}{s}\right)^n (n-1)^{n+1} - n(n-2)(n-1)^{n-1}\left(\frac{a}{s}\right)^n}{nv^{n-1}} \\
b(v) &= \frac{n-1}{n}v + \frac{a}{s} - \left(\frac{a}{s}\right)^n (n-1)^{n-1} \frac{(n-1)^2 - n(n-2)}{nv^{n-1}} \\
b(v) &= \frac{n-1}{n}v + \frac{a}{s} - \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{nv^{n-1}}
\end{aligned}$$

Next, we show that this is an equilibrium that maximizes utility. We assume a bidder can bid like a bidder with type x while every other bidder follows the derived bidding strategy. We show that it is optimal for a bidder not to deviate ($x = v$).

The expected utility of a bidder is $u(x, v) = x^{n-1}(a + s(v - b(x)))$

$$\begin{aligned}
\max_b u &: x^{n-1} \left(a + s \left(v - \frac{n-1}{n}x - \frac{a}{s} + \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{nx^{n-1}} \right) \right) \\
&= x^{n-1}sv - x^n s \frac{n-1}{n} + s \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{n}
\end{aligned}$$

We take the derivative to solve for the optimal x .

$$\begin{aligned}\frac{\partial u}{\partial x} &= (n-1)x^{n-2}sv - nx^{n-1}s\frac{n-1}{n} = 0 \\ (n-1)x^{n-2}sv &= nx^{n-1}s\frac{n-1}{n} \\ v &= x\end{aligned}$$

We take the second derivative to show that this is indeed a maximum.

$$\frac{\partial^2 u}{\partial x^2} = (n-2)(n-1)svx^{n-3} - (n-1)^2sx^{n-2}$$

This is a maximum because the function is concave around $v = x$.

$$(n-2)(n-1)svx^{n-3} - (n-1)^2sx^{n-2} < 0$$

$$n-2 < n-1 \quad \square$$

In the next step, we show that for values higher than $v > \bar{v} = \left(\frac{a}{s}\right)(n-1)$ the optimal bid according to the bidding function $b(v) = \frac{n-1}{n}v + \frac{a}{s} - \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{nv^{n-1}}$ is never higher than v . If the bid would be greater than v the manager would again restrain himself from bidding higher than v due to possible punishment from the owner. Hence, showing that the slope of the bidding function is never greater than 1 is a proof that the bidding function $b(v)$ is never higher than v for values $v > \bar{v}$.

$$\begin{aligned}b(v) &= \frac{n-1}{n}v + \frac{a}{s} - \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{nv^{n-1}} \\ &= \frac{n-1}{n}v + \frac{a}{s} - \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{n}v^{1-n}\end{aligned}$$

$$\begin{aligned} \frac{\partial b(v)}{\partial v} &< 1 \\ \frac{n-1}{n} + \left(\frac{a}{s}\right)^n \frac{(n-1)^n}{n} \frac{1}{v^n} &< 1 \\ n-1 + \left(\frac{a}{s}\right)^n (n-1)^n \frac{1}{v^n} &< n \\ \left(\frac{a}{s}\right)^n (n-1)^n \frac{1}{v^n} &< 1 \\ \left(\frac{a}{s}\right)^n (n-1)^n &< v^n \\ \frac{a}{s}(n-1) &< v \end{aligned}$$

Given the threshold \bar{v} the value v is greater than $\frac{a}{s}(n-1)$. Hence, the slope of the bidding function is lower than 1.

A.3 Additional Proofs

A.3.1 Manager's Bid greater than Bid in Standard Setting.

The following proof shows that the bid of a managers is always higher if $a > 0$ than compared to the standard setting.

Proof. Proof for values $v \leq \bar{v}$ that $b_{Standard}^{FPSB} < b_{Punishment}^{FPSB}$:

$$\begin{aligned} \frac{n-1}{n}v &< v \\ \frac{n-1}{n} &< 1 \end{aligned}$$

Proof for values $v > \bar{v}$ that $b_{Standard}^{FPSB} < b_{Punishment}^{FPSB}$:

$$\begin{aligned}
b_{Punishment}^{FPSB} &> b_{Standard}^{FPSB} \\
\frac{n-1}{n}v + \frac{a}{s} - \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{nv^{n-1}} &> \frac{n-1}{n}v \\
\frac{a}{s} &> \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{nv^{n-1}} \\
nv^{n-1} &> \left((n-1)\frac{a}{s}\right)^{n-1} \\
n^{(\frac{1}{n-1})}v &> (n-1)\frac{a}{s}
\end{aligned}$$

Only managers with a valuation v greater than the threshold $\bar{v} = \frac{a}{s}(n-1)$ bid according to the bidding function ($b(v) = \frac{n-1}{n}v + \frac{a}{s} - \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{nv^{n-1}}$). It is enough to prove that the condition holds for \bar{v} because higher values of v increase the left hand side of the equation. If the equation holds for $\bar{v} = \frac{a}{s}(n-1)$ it must hold for all values greater than \bar{v} .

$$\begin{aligned}
n^{(\frac{1}{n-1})}\frac{a}{s}(n-1) &> (n-1)\frac{a}{s} \\
n^{(\frac{1}{n-1})} &> 1 \\
n &> 1^{n-1}
\end{aligned}$$

By assumption $n > 1$. The condition holds. \square

A.3.2 Monotonic Increasing Bidding Function

This proof shows that the bidding function is monotonic increasing.

Proof.

$$\begin{aligned}
\frac{\partial b(v)}{\partial v} &= \frac{n-1}{n} - \left(\frac{a}{s}\right)^n \frac{(n-1)^{n-1}}{n} (1-n)v^{-n} \\
&= \frac{n-1}{n} + \left(\frac{a}{s}\right)^n \frac{(n-1)^n}{n} v^{-n} \\
&= \underbrace{\frac{n-1}{n}}_{>0} + \underbrace{\left(\frac{a}{s}\right)^n \frac{(n-1)^n}{n} \frac{1}{v^n}}_{>0} > 0
\end{aligned}$$

\square

Appendix B

Appendix to Chapter 2

B.1 Behavioral Theory Model

B.1.1 Proof of Proposition 2.1

Proof. By Assumptions 2.1 and 2.2, an increase in l_r increases a_p and c_r . We use the Implicit Function Theorem to prove that an increase in a_p or c_r both weakly increases c_p . Hence, the increase in l_r must weakly increase c_p . We start with Equation (2.2) derived in Section 2.4.1, which shows the condition that maximizes the adjusted utility of a proposer, assuming an interior solution.

$$F(c_p; a_p, c_r) = \frac{\partial v_p}{\partial c_p} = \underbrace{\frac{\partial u_p}{\partial c_p}}_{<0} + a_p \cdot \underbrace{\frac{\partial u_r}{\partial c_p}}_{>0} = 0$$

We can make statements about the first and second partial derivatives of the twice differentiable concave direct utility functions ($u_{p,r}$). Higher contributions by the proposer c_p increase the monetary outcome of a receiver and decrease the monetary outcome of a proposer. This means that a higher contribution by the proposer (c_p) has a negative effect on the proposer's direct utility, while it positively affects the receiver's direct utility ($\frac{\partial u_p}{\partial c_p} < 0$, $\frac{\partial u_r}{\partial c_p} > 0$). The second partial derivatives, $\frac{\partial^2 u_p}{\partial c_p^2} < 0$ and $\frac{\partial^2 u_r}{\partial c_p^2} < 0$, are both negative. The positive marginal utility of more money decreases for the receiver. The negative marginal

utility of losing money increases with less money for the proposer. The mixed partial derivatives are both positive ($\frac{\partial^2 u_p}{\partial c_p \partial c_r} > 0$ and $\frac{\partial^2 u_r}{\partial c_p \partial c_r} > 0$). For higher contributions of the other player, the negative marginal utility of contributing to the PGG is smaller, because the income is higher. This is true for the proposer and the receiver.

We use the Implicit Function Theorem to show how a change of a_p and c_r affects c_p . Proof that the optimal contribution c_p increases with a higher level of altruism ($\frac{\partial c_p}{\partial a_p} > 0$):

$$\frac{\partial c_p}{\partial a_p} = -\frac{\frac{\partial F}{\partial a_p}}{\frac{\partial F}{\partial c_p}} = -\frac{\overbrace{\frac{\partial u_r}{\partial c_p}}^{>0}}{\underbrace{\frac{\partial^2 u_p}{\partial c_p^2}}_{<0} + \underbrace{a_p \frac{\partial^2 u_r}{\partial c_p^2}}_{<0}} > 0$$

Proof that the optimal contribution c_p increases with a higher contribution of the receiver ($\frac{\partial c_p}{\partial c_r} > 0$):

$$\frac{\partial c_p}{\partial c_r} = -\frac{\frac{\partial F}{\partial c_r}}{\frac{\partial F}{\partial c_p}} = -\frac{\overbrace{\frac{\partial^2 u_p}{\partial c_p \partial c_r}}^{>0} + \overbrace{a_p \frac{\partial^2 u_r}{\partial c_p \partial c_r}}^{>0}}{\underbrace{\frac{\partial^2 u_p}{\partial c_p^2}}_{<0} + \underbrace{a_p \frac{\partial^2 u_r}{\partial c_p^2}}_{<0}} > 0$$

The equations above show that the denominator $\partial F/\partial c_p$ is always smaller than 0. Therefore, the necessary condition for the Implicit Function Theorem holds that the denominator is never 0.

This proves that c_p increases in l_r in the case of interior solutions. If the level of altruism a_p is so low that the contribution before and after the update is equal to zero ($\bar{c}_p = \check{c}_p = 0$), or if the contribution before is already at a maximum $\bar{c}_p = c^{max}$, the effect can be zero. Hence, the overall effect of an increase in l_r is non-negative on c_p . \square

B.1.2 Proof of Proposition 2.2


Proof. To prove that the proposer's adjusted utility (v_p) increases in l_r , we show that a proposer can always choose a contribution \check{c}_p that guarantees him a higher adjusted utility (v_p) than with a lower l_r . Following the experimental framework, we model an increase in l_r through learning the preferences of the matched receiver. This means that we demonstrate that a proposer's adjusted utility increases when he learns that l_r is higher than he previously thought. Note that we do not derive a proposer's optimal strategy, but show that there is always a strategy that makes the proposer better off.

The initially optimal contributions (given \bar{l}_r and \bar{a}_p) by a proposer (receiver) are denoted by \bar{c}_p (\bar{c}_r). The resulting monetary outcome of a proposer (receiver) is \bar{m}_p (\bar{m}_r) and their direct (monetary) utility is \bar{u}_p (\bar{u}_r). The preferences that the proposer then learns are denoted as \check{l}_r ($> \bar{l}_r$), and the receiver's contribution is \check{c}_r ($> \bar{c}_r$). The latter directly follows from Assumption 2. Note that if a player contributes c to the PGG, the sum of marginal returns for both players is greater than c . Therefore, contributing is always socially optimal.

In order to guarantee a higher adjusted utility, the proposer follows the following strategy: Contribute \check{c}_p , so that the receiver's new monetary outcome \check{m}_r equals her old monetary outcome \bar{m}_r (see Case 1). If this is not possible because it would require a higher contribution than is possible in the PGG ($\check{c}_p > c^{max}$), contribute the maximum possible contribution c^{max} to the PGG (see Case 2).

Case 1: Contribute \check{c}_p such that $\check{m}_r = \bar{m}_r$.

The receiver's direct utility \check{u}_r remains the same as her previous direct utility \bar{u}_r . Because both players contribute more, the overall monetary outcome is larger than before. Given that \check{c}_p is set such that $\check{u}_r = \bar{u}_r$, the monetary payoff for proposer (\check{m}_p) must have increased. This implies that the proposer's adjusted

Figure B.1: Contributing c^{max} by the Proposer

utility must also increase, because his direct utility u_p and the level of altruism a_p increases, while the receiver's direct utility remains constant u_r .

This strategy might not always be possible. It can be the case that, even if the proposer contributes c^{max} , the new receiver's monetary outcome remains smaller than before ($\check{m}_r < \bar{m}_r$). Nevertheless, contributing c^{max} will always yield a higher adjusted utility for the proposer v_p than before.

Case 2: Contribute $\check{c}_p = c^{max}$.

If the proposer contributes c^{max} and $\check{m}_r < \bar{m}_r$, the overall monetary outcome increases due to the increased overall contributions ($\check{m}_p + \check{m}_r > \bar{m}_p + \bar{m}_r$). Since $\check{m}_r < \bar{m}_r$, the monetary gain for the proposer must be greater than the monetary loss for the receiver ($\check{m}_p - \bar{m}_p > \bar{m}_r - \check{m}_r$). It must also follow that $\check{m}_r \geq \check{m}_p$ because the proposer contributes c^{max} . Only, if the receiver also contributes $\check{c}_r = c^{max}$, both monetary outcomes are the same ($\check{m}_p = \check{m}_r$). Due to the concavity of the direct utility function, the increase in proposer's direct utility must be greater than the direct utility loss for the receiver (see Figure B.1). The increase of altruism even dampens the decrease of the receiver's direct utility u_r on the proposer's adjusted utility v_p . \square

B.2 Preregistered Analyses

Result 2: Regression Analysis

Table B.1: Preference Adjustments across Treatments

	1[Preference Adjustment]	
	(1)	(2)
Info	.165*** [.095,.234]	.158*** [.089,.227]
Loss Aversion		-.020 [-.051,.012]
Cognitive Ability (Raven's')		-.007 [-.037,.022]
Male		-.002 [-.073,.070]
Observations	575	575

Notes. Logit Regressions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The table shows marginal effects at the mean from a logit regression. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals.

Result 3: Regression Analysis

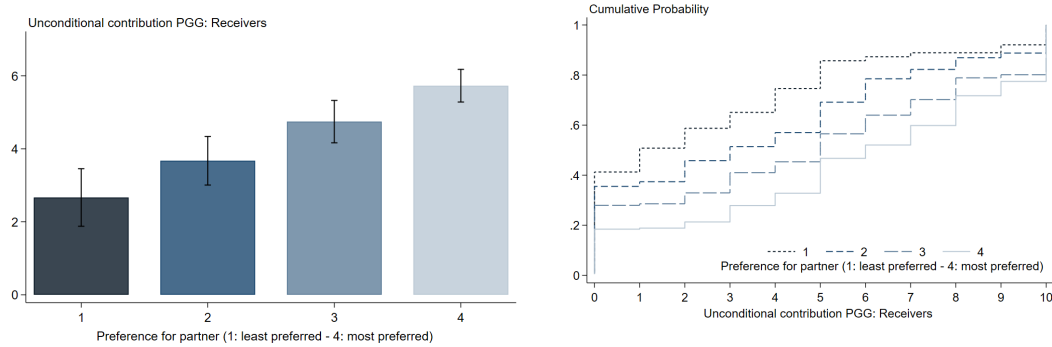
Table B.2: Consistency of Preference Adjustments

	1[Consistent Preference Adjustment]	
	(1)	(2)
Info	.152*** [.108,.195]	.150*** [.107,.193]
Loss Aversion		-.010 [-.024,.005]
Cognitive Ability (Raven's')		-.000 [-.016,.016]
Male		-.021 [-.060,.018]
Observations	575	575

Notes. Logit Regressions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. This table shows marginal effects at the mean from logit regressions where the dependent variable is an indicator for whether someone changed their preferences consistent with having reciprocal preferences. Standard errors are clustered at the individual level. The values in square brackets represent the 95% confidence intervals.

Info is an indicator, taking the value of one if the participant was randomly assigned to the treatment *Info*. *Loss aversion* and *Cognitive ability* are calculated as detailed in Footnote 17, *Male* is an indicator taking the value of 1 if a participant indicated to identify as male.

Unconditional Contributions of Receivers



A) Averages by *Preference for partner*

B) Distributions by *Preference for partner*

Notes. This figure displays the unconditional contributions of receivers by their preferences for the matched proposer. *Preference for partner (1-4)* takes the value of four if the matched proposer was the first choice of the receiver, three if the matched receiver was the second choice, and so on. Panel A) shows averages, Panel B) the cumulative distribution functions.

Figure B.2: Unconditional PGG Contributions: Receiver

Table B.3: Unconditional PGG Contributions of Receivers

	Unconditional PGG Contribution (0-10)	
	(1)	(2)
Preference for partner (1-4)	1.023*** [.719,1.328]	.960*** [.666,1.253]
Round		-.216*** [-.345,-.087]
Loss Aversion		-.480* [-.987,.027]
Cognitive Ability (Raven's')		.315 [-.090,.720]
Male		-.805 [-2.207,.597]
Observations	575	575

Notes. OLS Regressions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals.

Preference for partner (1-4) takes the value of four if the matched partner was the first choice of the participant, three if the matched partner was the second choice, and so on. *Round* is a count variable, indicating the number of the current round (Round 1-5). *Loss aversion* and *Cognitive ability* are calculated as detailed in Footnote 17, *Male* is an indicator taking the value of 1 if a participant indicated to identify as male.

B.3 Exploratory Analyses

Determinants of Proposers' Preference Adjustments

Table B.4: Determinants of Proposers' Preference Adjustments

	1[Preference Adjustment]		
	No-Info	Info	
	(1)	(2)	(3)
Preference for initial partner (1-4)	-.122*** [-.172,-.073]	-.117*** [-.180,-.055]	-.124*** [-.188,-.060]
Round	-.023** [-.045,-.001]	-.027* [-.056,.003]	-.028* [-.058,.002]
Loss Aversion	-.004 [-.047,.040]	-.024 [-.081,.033]	-.025 [-.082,.033]
Cognitive Ability (Raven's')	.001 [-.025,.028]	-.020 [-.080,.041]	-.022 [-.084,.039]
Male	.036 [-.055,.127]	-.047 [-.165,.071]	-.046 [-.164,.073]
Average preference of other receivers (1-4)		.088** [.010,.166]	
Initial partner's preference (1-4)		-.081*** [-.140,-.021]	-.085*** [-.145,-.025]
Highest preference of other receivers (1-4)			.092*** [.041,.143]
Observations	290	285	285

Notes. OLS Regressions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals.

Preference for initial partner (1-4) takes the value of four if the initial matched partner was the first choice of the participant, three if the matched partner was the second choice, and so on. *Round* is a count variable, indicating the number of the current round (Round 1-5). *Loss aversion* and *Cognitive ability* are calculated as detailed in Footnote 17, *Male* is an indicator taking the value of 1 if a participant indicated to identify as male. *Initial partner's preferences (1-4)* takes the value of four if the participant was the most preferred choice of their initial partner (i.e. before being able to adjust their preferences), three if the participant was the second most preferred choice, and so on. *Average preference of other receivers (1-4)* calculates the average preference of the other receiver and takes the value of four if the participant was the most preferred choice of all three receivers, the participant was not matched to initially. *Highest preference of other receivers (1-4)* takes the value of four if the partner was the most preferred choice of at least one of the non-matched receivers, three if the participant was not the most preferred choice of any receiver, but the second most preferred choice of at least one, and so on.

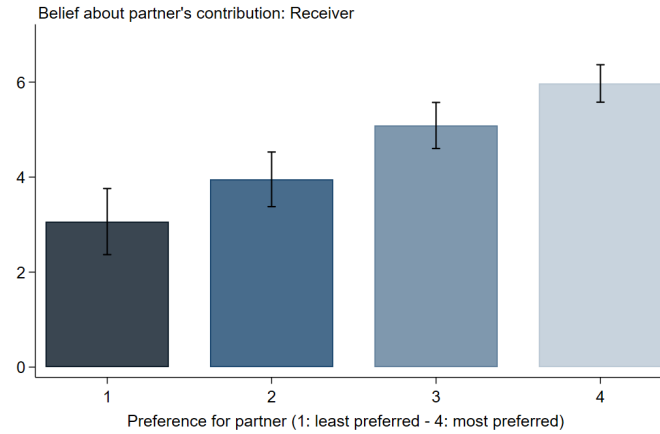
Predicting PGG Contribution with Questionnaire Responses

Table B.5: PGG Contributions and Questionnaire Responses

	Unconditional Contribution	Avg. Conditional Contribution
	(1)	(2)
Cat over Dog	-.207 [-.535,.121]	.106 [-.204,.415]
Book over Film	.372 [-.119,.863]	.005 [-.396,.406]
Beach over City	.150 [-.343,.644]	-.028 [-.410,.353]
Bar over Club	-.176 [-.678,.326]	-.225 [-.643,.192]
Living Alone over Shared	-.133 [-.531,.264]	-.117 [-.487,.253]
Reserved	.455* [-.025,.935]	.179 [-.249,.607]
Lazy	.014 [-.509,.537]	.021 [-.428,.470]
Handy with Hands	.261 [-.176,.698]	.257 [-.105,.619]
Spontaneous	.092 [-.421,.605]	.229 [-.290,.748]
Conflict Avoidant	.046 [-.456,.547]	.227 [-.175,.629]
Strictness Covid19 Policy	-.108 [-.691,.475]	.333 [-.097,.763]
Quota Disadvantaged	.417 [-.081,.914]	-.031 [-.449,.387]
Bicycle Helmet Mandatory	.032 [-.420,.485]	.055 [-.318,.428]
Legalize Marijuana	.342 [-.092,.775]	.194 [-.219,.606]
Taxes Unhealthy Food	-.124 [-.543,.296]	.106 [-.229,.441]
Observations	1150	575

Notes. OLS Regressions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals. Column (1) includes both receivers and proposers. Column (2) only includes proposers, because receivers did not make conditional contribution decisions. For the wording of the questions, answered on a Likert scale from 1-4, see Appendix B.4.2.

Beliefs of Receivers about PGG Contribution of Partner



Notes. This figure displays the beliefs of receivers about the unconditional PGG contributions of their matched proposer by their preferences for the matched proposer. *Preference for partner (1-4)* takes the value of four if the matched proposer was the first choice of the receiver, three if the matched receiver was the second choice, and so on.

Figure B.3: Beliefs of Receivers: PGG Contributions of Partner

Table B.6: Beliefs of Receivers: PGG Contributions of Partner

	Beliefs about partner's PGG contribution (0-10)	
	(1)	(2)
Preference for partner (1-4)	.983*** [.742,1.225]	.944*** [.708,1.181]
Round		-.057 [-.198,.084]
Loss Aversion		-.234 [-.710,.241]
Cognitive Ability (Raven's')		.213 [-.111,.537]
Male		-.056 [-1.142,1.029]
Observations	575	575

Notes. OLS Regressions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals.

Preference for partner (1-4) takes the value of four if the matched partner was the first choice of the participant, three if the matched partner was the second choice, and so on. *Round* is a count variable, indicating the number of the current round (Round 1-5). *Loss aversion* and *Cognitive ability* are calculated as detailed in Footnote 17, *Male* is an indicator taking the value of 1 if a participant indicated to identify as male.

Payoffs from PGG across Treatments




Figure B.4: Payoffs PGG: Implementation of Unconditional Decisions

B.4 Instructions

Appendix B.4 includes the translated instructions of the experiment (from German). Treatment specific parts are shown in *italics* and the corresponding treatment is clearly indicated.

B.4.1 General Instructions (before Part I)

**Welcome to the experiment and thank you for
your participation!**

Please do not speak from now on with any other participant.

Procedures

In this experiment, we study economic decision-making. You can earn money by participating. The money you earn will be paid to you privately after the experiment.

The experiment lasts around 90 minutes and consists of four parts (I-IV). At the beginning of every part, you receive detailed instructions. In addition, you will receive comprehension questions for some parts to help you understand how the experiment works and the payoff conditions. If you have questions after reading the instructions or during the experiment, please raise your hand or press the red button on your keyboard. One of the experimenters will then come to you and answer your questions privately.

Tools

You find a pen at your desk. Please leave the pen and the instructions on the table after the experiment.

Anonymity

The analysis of the experiment is anonymous; that is, we will never link your name with the data generated in the experiment. To receive your payoff, you will need to provide your bank details or PayPal mail address at the end of the experiment. No further personal data will be passed on. Information collected during the experiment may be visible to other participants as the experiment progresses. You make all decisions anonymously, so no other participant can associate your decisions with you during the experiment.

Payment

In addition to the income that you earn during the experiment, you will receive 6 € for showing up on time and answering a short questionnaire. In addition, you can achieve additional payoffs during the experiment. During the experiment, you and the other participants will be asked to make a series of decisions. These can affect the payoffs for you, and potentially for other participants. Additionally, you can earn money by making correct assessments. How your decisions relate to the payoffs will be explained in more detail in the respective instructions.

Exchange rate

In some parts of the experiment, we do not talk about Euro, but about Taler.

We convert Taler into Euros at the end of the experiment. Please note the following exchange rate:

$$1 \text{ Taler} = 0,70 \text{ €}$$

B.4.2 Questionnaire (Part I)

[Instructions: *In the first part of the experiment, we ask you to truthfully fill out a questionnaire. This is a personality questionnaire, so there are no right or wrong answers.*

Please answer the questions with the answer options:

- *Does not apply* • *Tends not to apply* • *Tends to apply* • *Applies*]

1. I would rather have a cat than a dog as a pet.
2. I prefer reading a book in the evening to watching a movie.
3. I prefer to go to the beach on vacation than to visit a city.
4. I would rather spend an evening in a bar than partying in a club.
5. I prefer to live in a shared apartment than alone.
6. I am rather reserved and quiet.
7. I am easygoing, prone to laziness.
8. I am talented with my hands.
9. I often make decisions spontaneously and intuitively.
10. I tend to avoid conflict.
11. I am in favor of strong policy measures to contain the Covid-19 pandemic in Germany.
12. I support quota regulations in the labor market for socially disadvantaged groups (e.g., for women or migrants).
13. There should be a requirement to wear a bicycle helmet.
14. The possession of marijuana should be legalized.
15. Unhealthy foods should be taxed more.

B.4.3 Instructions (Part II)

The participants received the instructions for Part II of the experiment in print. An interactive screen to familiarize with the matching procedure and control questions to ensure understanding were later displayed on the computer screens.

Proposer

Part II of the experiment consists of 5 rounds. Each round is structured in the same way. In each round, you will make decisions that affect your payout amount, as well as the payout amount of another participant. One round will be randomly selected for which the achieved amount will be paid out. You will find out which round was selected only at the end of the experiment. Therefore, you should carefully consider your decisions in all rounds, as each may become relevant to you.

You were randomly assigned one of two roles for Part II of the experiment. This role remains the same across all rounds. There are participants of "Type P" and participants of "Type R". You are "Type P". All participants of "Type P" receive identical instructions. Participants of "Type R" are in a similar decision situation, we explicitly point out any differences. In each round, four "Type P" participants are matched with four "Type R" participants. This means that 8 randomly selected participants interact with each other per round. In each round, you will be randomly selected to interact with other participants.

We will illustrate the process of Part II using one round as an example. We will refer to your group of four "Type P" participants as Group A, and to the group of four "Type R" participants with whom you interact as Group B.

Each round consists of three consecutive sections (Section 1, Section 2 and Section 3).

In the final Section 3, you will simultaneously make decisions with one participant from Group B (your team partner) that are payoff-relevant for both of you.

In Section 3, one participant from Group A and one participant from Group B thus form a team of 2.

In Section 1, you specify which participant of Group B you want as your team partner in this decision situation. Your choice of team partner is important to you because your team partner's decisions affect your payoffs.

In Section 2, you will be assigned a team partner for Section 3 based on your choice and the choices of the other participants through an assignment mechanism.

Below you find detailed information on all three sections.

Section 1

In the first section, you will see a randomly selected part of the answers of the 4 participants of Group B from the questionnaire. These participants are your possible team partners.

Example image: Answers from the questionnaire

	Participant A	Participant B	Participant C	Participant D
Statement 1: I support quota regulations in the labor market for socially disadvantaged groups (e.g., for women or migrants).	Applies	Tends to apply	Tends not to apply	Tends not to apply
Statement 2: There should be a requirement to wear a bicycle helmet.	Does not apply	Applies	Does not apply	Does not apply
Statement 3: I would rather spend an evening in a bar than partying in a club.	Applies	Tends to apply	Does not apply	Applies
Statement 4: I am easygoing, prone to laziness.	Tends to apply	Tends to apply	Tends to apply	Does not apply
Statement 5: I prefer to live in a shared apartment than alone.	Does not apply	Does not apply	Tends to apply	Tends to apply

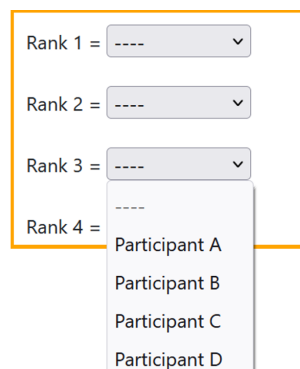
Figure B.5: Instruction - Proposer: Questionnaire

At the same time, the participants of Group B (Participants A-D) see other randomly selected answers from your questionnaire and the questionnaires of the other 3 participants of Group A.

After viewing the profiles, we ask you to submit a preference order.

With this preference order, you indicate with whom of the participants from Group B you would prefer to be in the decision situation in Section 3. Rank 1 means that you would most like to have this participant as your team partner. Rank 2 means that you would second most like to have this participant as your team partner, and so on.

Example image: Preference order



The image shows a form for submitting a preference order. It consists of four rows, each with a label 'Rank 1 =', 'Rank 2 =', 'Rank 3 =', and 'Rank 4 =' followed by a dropdown menu. The dropdown menus are currently empty, showing only a downward arrow. The 'Rank 4 =' dropdown menu is open, showing a list of four options: 'Participant A', 'Participant B', 'Participant C', and 'Participant D'. The entire form is enclosed in a thin orange border.

Figure B.6: Instruction - Proposer: Preference Order

All other participants of Groups A and B will also be asked to submit such a preference order.

Section 2

In this section, a two-step mechanism will determine the allocation for Section 3. The mechanism is chosen so that it is always best for you to submit your actual preference order.

Example: Suppose you could choose between participants A, B, C or D from Group B. If you would prefer to have Participant A, second favorite Participant B, third favorite Participant C, and fourth favorite Participant D as your team partner, then you should submit the preference order $A > B > C > D$. If the assignment mechanism assigned you Participant B, for example, under the submission of your true preference order, there is no other preference order by which the mechanism assigns Participant A to you.

Your preference order

Your preference order	Rank on which the respective participant has placed you.
Rank 1: Participant A	Rang 3
Rank 2: Participant B	Rang 2
Rank 3: Participant C	Rang 1
Rank 4: Participant D	Rang 3

Your assigned team partner: Participant A

Below you can find the information from Part I about your potential team partners.

Information about your potential team partners

If you wish, you can adjust your preference order at this point. An adjustment makes sense if your preference order is different from the one you submitted previously (and which is displayed in the table on the left).

[Adjust preference order](#) [Continue](#)

No-Info

Figure B.7: Instruction - Proposer: Adjustment of preferences

In the first step, the allocation mechanism determines the 2-person teams based on the preferences submitted. Then you will see which participant of Group B has been assigned to you. *In addition, for each participant of Group B, you will see the rank they have placed you on. [Only in Info]*

Example screen: Adjustment of preferences

In this example, in the first step, you have set Participant A to Rank 1, and have been assigned him or her as a team partner by the mechanism. *Participant A has placed you on Rank 3 of their preference order. [Only in Info]*

If you wish, you can adjust your preference order at this point. An adjustment makes sense if your preference order is different from the one you submitted previously.

In the second step, the allocation mechanism again determines 2-person teams based on these preference orders. If at least one participant has adjusted their preference order, other teams may result compared to the teams after the first step. The key is that it is always best for you to submit your true preference order.

At the end of Section 2, it will be randomly selected whether your final team partner for Section 3 will be the one assigned to you after the first step, or whether your team partner will be the one assigned to you after the second step of Section 2. Therefore, you should submit your true preference order in both

steps.

Information and procedure for participants of Group B

The process of Section 2 is different for participants from Group B. Unlike you, your potential team partners from Group B cannot adjust their preference order in the second part of the assignment mechanism. *Participants from Group B do not know the preference orders of Group A and do not know that Group A will receive the preference order of Group B. [Only in Info]*

Section 3

Decision situation

You and your team partner can each put 10 Taler into a private account, or you can put all or part of 10 Taler into a joint account. Any money that you do not deposit into the joint account will automatically be deposited into the private account. You and your team partner will make your decisions independently and secretly in this part.

Income from the private account

Every Taler you put on the private account, you will get paid at the end. If you keep 10 Taler for yourself, you will receive these 10 Taler from the private account. If you keep 6 Taler for yourself, you will receive these 6 Taler from the private account. Nobody but you receives income from your private account.

Income from the joint account

You can also put your Taler into the joint account. For each Taler contributed to the joint account, both you and your team partner will receive 0.75 Taler each. Both of you benefit from the joint account to the same extent, regardless of your respective deposits. The payoff from the joint account depends only on the sum of the deposits.

The payout of each team member is determined by the following formula.

$$\text{Individual payout for each team member} = \frac{(\text{deposit from you} + \text{deposit from your team partner}) * 0.75}{2}$$

If you and your team partner deposit 5 Taler each, the sum of the two deposits is $5+5=10$. Of these 10 Taler, you and your team partner will each receive $10 \cdot 0.75 = 7.5$ Taler. If you and your team partner deposit a total of 16 Taler, you will both receive $16 \cdot 0.75 = 12$ Taler.

Total income

Your total income is the sum of your income from the personal account and your income from the joint account.

Your input

You and your team partner from Group B simultaneously and independently make the decision how many of your 10 Taler you want to contribute to the joint account. We call this decision contribution in the following.

In addition to this, participants in Group A make a second contribution decision, the contribution table. For participants of Group A, it is chosen at random whether the contribution or the contribution table is relevant for payout. You must therefore carefully consider both types of contribution decisions, as both may become relevant to you. Since participants of Group B only make the contribution decision, the contribution is always and exclusively payoff relevant for these participants.

Contribution and contribution table

With your contribution to the joint account, you determine how many of the 10 Taler you want to deposit into the joint account. The deposit to your private account is automatically the difference between 10 Taler and your contribution to the joint account.

Example image: Contribution

Please indicate the amount you wish to deposit into the joint account:

	0 1 2 3 4 5 6 7 8 9 10
How many of your 10 Taler do you contribute to the joint account?	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/>

Figure B.8: Instruction - Proposer: Contribution

In the contribution table, you specify how many Taler you want to contribute to the joint account for each possible contribution of your team partner. So you make your own contribution decision based on how much your team partner contributes.

Example image: Contribution table

For each possible contribution of your team partner, please indicate the amount you would like to contribute to the joint account (of course, you can choose the same amount more than once):

How many of your 10 Taler do you contribute to the joint account if....

	0	1	2	3	4	5	6	7	8	9	10
...Your team partner contributes 0 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
...Your team partner contributes 1 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
...Your team partner contributes 2 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
...Your team partner contributes 3 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
...Your team partner contributes 4 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
...Your team partner contributes 5 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
...Your team partner contributes 6 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
...Your team partner contributes 7 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
...Your team partner contributes 8 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
...Your team partner contributes 9 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
...Your team partner contributes 10 Taler?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure B.9: Instruction - Proposer: Contribution table

After the decision:

You will find out the result of the selected round only at the end of the experiment.

You can now familiarize yourself on the computer monitor with both the submission of preference sequences, as well as the allocation mechanism. After that, you will get some comprehension questions.

Receiver

Part II of the experiment consists of 5 rounds. Each round is structured in the same way. In each round, you will make decisions that affect your payout amount, as well as the payout amount of another participant. One round will be randomly selected for which the achieved amount will be paid out. You will find out which round was selected only at the end of the experiment. Therefore, you should carefully consider your decisions in all rounds, as each may become relevant to you.

You were randomly assigned one of two roles for Part II of the experiment. This role remains the same across all rounds. There are participants of "Type P" and participants of "Type R". You are "Type R". All participants of "Type R" receive identical instructions. Participants of "Type P" are in a similar decision situation. In each round, four "Type P" participants are matched with four "Type R" participants. This means that 8 randomly selected participants interact with each other per round. In each round, you will be randomly selected to interact with other participants.

We will illustrate the process of Part II using one round as an example. We will refer to your group of four "Type P" participants with whom you interact as Group A, and your group of four "Type R" participants as Group B.

Each round consists of three consecutive sections (Section 1, Section 2 and Section 3).

In the final Section 3, you will simultaneously make decisions with one participant from Group A (your team partner) that are payoff-relevant for both of you. In Section 3, one participant from Group A and one participant from Group B thus form a team of 2.

In Section 1, you specify which participant of Group A you want as your team partner in this decision situation. Your choice of team partner is important to you because your team partner's decisions affect your payoffs.

In Section 2, you will be assigned a team partner for Section 3 based on your choice and the choices of the other participants through an assignment mechanism.

Below you find detailed information on all three sections.

Section 1

In the first section, you will see a randomly selected part of the answers of the 4 participants of Group A from the questionnaire. These participants are your possible team partners.

Example image: Answers from the questionnaire

	Participant A	Participant B	Participant C	Participant D
Statement 1: I support quota regulations in the labor market for socially disadvantaged groups (e.g., for women or migrants).	Applies	Tends to apply	Tends not to apply	Tends not to apply
Statement 2: There should be a requirement to wear a bicycle helmet.	Does not apply	Applies	Does not apply	Does not apply
Statement 3: I would rather spend an evening in a bar than partying in a club.	Applies	Tends to apply	Does not apply	Applies
Statement 4: I am easygoing, prone to laziness.	Tends to apply	Tends to apply	Tends to apply	Does not apply
Statement 5: I prefer to live in a shared apartment than alone.	Does not apply	Does not apply	Tends to apply	Tends to apply

Figure B.10: Instruction - Receiver: Questionnaire

At the same time, the participants of Group A (Participants A-D) see other randomly selected answers from your questionnaire and the questionnaires of the other 3 participants of Group B.

After viewing the profiles, we ask you to submit a preference order.

With this preference order, you indicate with whom of the participants from Group A you would prefer to be in the decision situation in Section 3. Rank 1 means that you would most like to have this participant as your team partner.

Rank 2 means that you would second most like to have this participant as your team partner, and so on.

Example image: Preference order

Figure B.11: Instruction - Receiver: Preference order

All other participants of Groups A and B will also be asked to submit such a preference order.

Section 2

In this section, a mechanism will determine the allocation for Section 3. The goal of the mechanism is to assign participants their best possible team partner. The mechanism is based on a simple logic: If several participants of Group A want you to be their team partner, the mechanism will always select for you the participant that you have specified further ahead in your preference order.

Example: Suppose you could choose between participants A, B, C or D from Group A. You prefer to have Participant A, second favorite Participant B, third favorite Participant C, and fourth favorite Participant D as your team partner ($A > B > C > D$). If the assignment mechanism does not assign you Participant A when you state your true preference order, it automatically means that Participant A prefers another participant of Group B over you.

Let us assume that this is the case. Now, if both participant B and C would prefer you to be their team partner, the mechanism will choose the participant you have specified further up in your preference order as your team partner. If

you would submit the preference order $A > B > C > D$, you would get Participant B as your team partner. If you would give the preference order $A > C > B > D$, you would get Participant C as your team partner. This also means that if you submit a preference order that does not match your true preference order, you may not get your best possible team partner.

Once you have submitted your preference order, you cannot change it.

Section 3

Decision situation

You and your team partner can each put 10 Taler into a private account, or you can put all or part of 10 Taler into a joint account. Any money that you do not deposit into the joint account will automatically be deposited into the private account. You and your team partner will make your decisions independently and secretly in this part.

Income from the private account

Every Taler you put on the private account, you will get paid at the end. If you keep 10 Taler for yourself, you will receive these 10 Taler from the private account. If you keep 6 Taler for yourself, you will receive these 6 Taler from the private account. Nobody but you receives income from your private account.

Income from the joint account

You can also put your Taler into the joint account. For each Taler contributed to the joint account, both you and your team partner will receive 0.75 Taler each. Both of you benefit from the joint account to the same extent, regardless of your respective deposits. The payoff from the joint account depends only on the sum of the deposits.

The payout of each team member is determined by the following formula.

$$\text{Individual payout for each team member} = \frac{(\text{deposit from you} + \text{deposit from your team partner}) * 0.75}{2}$$

If you and your team partner deposit 5 Taler each, the sum of the two deposits is $5+5=10$. Of these 10 Taler, you and your team partner will each receive $10*0.75 = 7.5$ Taler. If you and your team partner deposit a total of 16 Taler, you will both receive $16*0.75 = 12$ Taler.

Total income

Your total income is the sum of your income from the personal account and your income from the joint account.

Your input

You and your team partner from Group B simultaneously and independently make the decision how many of your 10 Taler you want to contribute to the joint account. We call this decision contribution in the following.

Contribution

With your contribution to the joint account, you determine how many of the 10 Taler you want to deposit into the joint account. The deposit to your private account is automatically the difference between 10 Taler and your contribution to the joint account.

Example image: Contribution

Please indicate the amount you wish to deposit into the joint account:

	0	1	2	3	4	5	6	7	8	9	10
How many of your 10 Taler do you contribute to the joint account?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure B.12: Instruction - Receiver: Contribution

After the decision:

You will find out the result of the selected round only at the end of the experiment.

You can now familiarize yourself on the computer monitor with both the submission of preference sequences, as well as the allocation mechanism. After that, you will get some comprehension questions.

B.4.4 Additional Instructions (Part III)

Beliefs

In this part of the experiment, we ask you to guess the decisions of your respective team partners from Part II. You thus provide an estimate for each of the rounds played. Your payoff depends on whether you estimate the contribution to the joint account of your respective team partner in Part II correctly.

Before each decision, you will again receive the information about your team partner that you had available when you made your own contribution decision. Please provide an estimate of how many Taler your respective team partner put into the joint account. *Note that your team partner made this decision, without knowing your submitted preference order. [only proposer]*

Payoff

If you estimate your team partner's contribution exactly correctly, you will receive 2 Euro for this correct estimation. If you estimate the contribution incorrectly, you will receive 0 Euro.

One of the rounds will be randomly selected for which the amount scored will be paid out. You will find out the result of the selected round only at the end of the experiment (after part IV).

Raven's Matrices

In this part of the experiment we ask you to complete figures. The figures consist of 3x3 elements that are logically connected. In each figure the lower right element is missing. We ask you to complete this with one of the 6 answer choices.

You have a total of 5 minutes to solve as many matrices as you can manage. The maximum number is 10 matrices. You will receive 0.50 Euro for each correctly solved matrix and 0.50 Euro will be deducted for each incorrectly solved matrix. You will receive at least 0.00 Euro for this task. You cannot get a negative payout from this task. Please select the appropriate image in each

case and confirm your selection. On the next page you can see an example.

Loss attitudes (Gächter, Johnson, & Herrmann, 2022)

This task consists of 6 decisions where you can accept up to 6 offers.

The offers consist of a lottery through which you can lose or win money. You have to decide for each of the 6 offers whether to accept it or not. For each accepted offer, the computer plays the lottery and hence decides if you lose or win money.

At the end of the experiment, your decision is implemented for one of the 6 offers. The computer randomly selects (with equal probability) which offer will be implemented.

Decide for each offer whether you want to accept it.

Table B.7: Instruction - Loss attitudes

1	With 50% probability you lose 2 Euro; with 50% probability you win 6 Euro.	<input type="radio"/> accept <input type="radio"/> reject
2	With 50% probability you lose 3 Euro; with 50% probability you win 6 Euro.	<input type="radio"/> accept <input type="radio"/> reject
3	With 50% probability you lose 4 Euro; with 50% probability you win 6 Euro.	<input type="radio"/> accept <input type="radio"/> reject
4	With 50% probability you lose 5 Euro; with 50% probability you win 6 Euro.	<input type="radio"/> accept <input type="radio"/> reject
5	With 50% probability you lose 6 Euro; with 50% probability you win 6 Euro.	<input type="radio"/> accept <input type="radio"/> reject
6	With 50% probability you lose 7 Euro; with 50% probability you win 6 Euro.	<input type="radio"/> accept <input type="radio"/> reject

Socio-demographics

Please provide the following statistical information.

- Age [integer]
- Gender [male; female; diverse]
- Field of study (faculty/major) [string]
- What language(s) is (are) your native language(s)? [string]
- What is your high school graduation grade? [number; 1-6]
- What is your high school graduation grade in mathematics? [number; 1-6]
- How many times have you participated in an economic laboratory study (including outside of this laboratory)? [0; 1-2; 3-5; 5+]

Appendix C

Appendix to Chapter 3

C.1 Proof of Proposition 3.1

In a matching market with reciprocal preferences, ex-post stability is a sufficient condition for immediate stability.

Proof. Assume a matching is immediate unstable but all the possible resulting outcomes are ex-post stable. A matching is immediate unstable if it is blocked by a pair or an individual after the mechanism took place, and every worker updates her belief about the types of the firms through seeing the final matching. When evaluating ex-post stability, workers learn the true type of their matched firm. Therefore, a change in stability between immediate and ex-post stability can only be attributed to workers receiving information about their matched partner's type. There is no information update for firms.

If a matching is immediate unstable because it is blocked by a pair, it must be the case that a worker (e.g. I) prefers another firm (e.g. B) over her matched firm (e.g. A). This means that the expected utility of being matched with firm B is higher than the expected utility of being matched with firm A . At the same time, if the matching is ex-post stable, she prefers to stay with her matched firm (A) over the other firm (B) for every possible realization of firm A 's type. It cannot be the case that being matched with any type of firm A is better than

the expected utility of being matched with firm B (ex-post stability), but the expected utility of being matched with firm A is lower than the expected utility of being matched with firm B (immediate instability).

The same logic applies for a worker that prefers being unmatched over her current match in terms of immediate stability. For a worker who is unmatched after the mechanism takes place, the information set is the same when evaluating immediate and ex-post stability. Hence, the matching cannot be immediate stable and ex-post unstable. \square

C.2 Proofs of Additional Statements

C.2.1 Statement 1

Truth-telling is not a weakly dominant strategy for firms with an anticipated rematching stage in a DA mechanism with reciprocal preferences.

Table C.1: Example 5

Proposer / Firm		Receiver / Worker
$A^1 : I \succ II$	with (p)	$I : A_1 \succ I \succ A_2$
$A^2 : III \succ I$	with $(1 - p)$	$II : B \succ A \succ C$
$B : III \succ II$		$III : C \succ A \succ B$
$C : II \succ III$		

Given:

$$I : u(A_1) \cdot p + u(A_2) \cdot (1 - p) > u(I)$$

Truth-telling is a weakly dominant strategy for proposers in a DA if receivers have reciprocal preferences (see Proposition 3.2). Example C.1 shows that this

is not the case if players can rematch after the matching and anticipate this. We show that a firm can improve its expected utility by deviating from truth-telling.

Proof. Truth-telling is a weakly dominant strategy if the outcome is never worse than any other strategy, given every possible strategy of all other players. We show that another strategy than truth-telling is better for firm A^2 given a rematching stage and given strategies of the other players. Assume that all proposers and workers I and II state true preferences. For receiver I , we assume that she states $A \succ I$. Independent of firm A 's type, worker I will always be matched with firm A , if both types of firm A state true preferences. However, the matching of the other firms and workers depends on the state preferences of firm A . In the event of A^1 , firm B is matched with worker III , and firm C is matched with worker II . If firm A is of type A^2 , firm B is matched with worker II and firm C is matched with worker III . After observing the matching, worker I infers the type of firm A and breaks up the matching if she knows that she is matched with firm A^2 . For this reason, firm A always states $I \succ II$. Worker I cannot tell which type she is matched with and will not break up the match with firm A . Hence, firm A^2 is mimicing the strategy of type A^1 and not stating true preferences. Because we are checking for a weakly dominant strategy, we do not have to check whether the played strategies are an equilibrium. \square

C.2.2 Statement 2

There are matching markets that are always ex-post unstable for any possible strategy a player can choose in the DA mechanism if all the other players state true preferences.

Table C.2: Example 6

Proposer / Firm		Receiver / Worker
$A^1 : I \succ II$	with (p)	$I : A_1 \succ I \succ A_2$
$A^2 : II \succ I$	with ($1 - p$)	$II : II$

Given:

$$I : u(I) < p \cdot u(A_1) + (1 - p) \cdot u(A_2)$$

Proof. The DA is not ex-post stable in our setting with reciprocal preferences. Example C.2 shows that no set of strategies always results in an ex-post stable matching. Firm A has two possible types, A^1 and A^2 and worker I has reciprocal preferences and would like to be matched when firm A ranks her first but wishes to be unmatched if the firm is of type A^2 . We assume that the expected utility from being matched with firm A without knowing its type is higher for worker I than being unmatched [$p \cdot u(A_1) + (1 - p) \cdot u(A_2) > u(I)$]. We show that both possibilities to submit preferences ($A \succ I$ or $I \succ A$) lead to ex-post unstable outcomes. If worker I decides to match with firm A (by stating $A \succ I$), the match is immediate stable. However, once she learns the type of firm A and it turns out to be of type A^2 , she prefers to be unmatched, and the matching is unstable. If worker I decides to remain unmatched initially (by stating $I \succ A$), she does not learn the type of firm A . However, the expected utility of being matched with firm A without knowing its type is higher than remaining unmatched. Firm A and worker I want to match, the matching is immediate and ex-post unstable. \square

C.2.3 Statement 3

There are markets where the DA mechanism leads to immediate instability, although some resulting outcomes are ex-post stable.

We use Example C.3 to show that there are markets in which the DA mechanism leads to immediate instability, although some outcomes are ex-post stable.

Table C.3: Example 7

Proposer / Firm		Receiver / Worker
$A^1 : I$	with (p)	$I : A_1 \succ B_1 \succ A_2 \succ B_2$
$A^2 : III \succ I$	with $(1 - p)$	$II : B$
$B^1 : I$	with (q)	$III : III$
$B^2 : III \succ I \succ II$	with $(1 - q)$	

Given:

$$I : p \cdot (A_1) + (1 - p) \cdot u(A_2) > q \cdot u(B_1) + (1 - q) \cdot u(B_2)$$

$$I : p \cdot (A_1) + (1 - p) \cdot u(A_2) < u(B_1)$$

Proof. In this example, every type of every firm will make an offer to worker I in a DA mechanism. Hence worker I states $A \succ B$ following her expected utility $[I : p \cdot (A_1) + (1 - p) \cdot u(A_2) > q \cdot u(B_1) + (1 - q) \cdot u(B_2)]$. After observing the matching, worker I can infer the type of firm B . If firm B is unmatched, it is of type B^1 , if firm B is matched to worker II , it is of type B^2 . When firm B is unmatched, the matching is immediate unstable because worker I prefers to be matched with firm B^1 over the lottery of being matched with either type A^1 or A^2 of firm A $[I : p \cdot (A_1) + (1 - p) \cdot u(A_2) < u(B_1)]$. However, if firm A is of type A^1 , worker I prefers to stay matched with firm A and the outcome is ex-post stable. \square

C.2.4 Statement 4

Proposers with reciprocal preferences do not have a weakly dominant strategy in the DA mechanism.

We show in Example C.4 that proposer I with reciprocal preferences does not have a weakly dominant strategy for preference reporting in a matching market when the DA is applied.

Table C.4: Example 8

Proposer / Worker	Receiver / Firm
$I : A_1 \succ B \succ A_2$	$A^1 : I$ with (p)
$II : A$	$A^2 : II \succ I$ with $(1 - p)$
	$B : I$

Given:

$$I : u(B) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$$

Proof. Given that proposer I prefers to be matched with firm B over being matched with firm A without knowing firm A 's type [$I : u(B) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$], the optimal strategy of applicant I depends on the strategy of other players. Here, the decision to state $A \succ B$ or $B \succ A$ depends on proposer II . If proposer II states her true preferences, proposer I states $A \succ B$ because she does not incur the risk be matched with A^2 . However, if applicant II states $II : II$, it is optimal for applicant I to state $I : B \succ A$ because she is matched with type A^2 with probability $(1 - p)$ if when stating $I : A \succ B$. This proves that there is no dominant strategy for applicant I . \square

C.2.5 Statement 5

The DA mechanism is not immediate and ex-post stable in the school choice setting, regardless of whether applicants are on the receiving or proposing side.

In Section 3.6, we already show in Example C.5 that the DA is neither immediate nor ex-post stable in a school choice setting, with applicants on the receiving side. This example shows an immediate and ex-post unstable matching market if applicants are on the proposing side.

Table C.5: Example 9

Proposer / Applicant	Receiver / School
$I : A_1 \succ B \succ A_2$	$A^1 : I$ with (p)
$II : II$	$A^2 : II \succ I \succ III$ with $(1 - p)$
$III : A$	$B : I$

Given:

$$I : u(B) > p \cdot u(A_1) + (1 - p) \cdot u(A_2)$$

Proof. Applicant I decides whether to state $A \succ B$ or $B \succ A$ and prefers being matched with school B over being matched with school A without an update about its type. Accordingly, applicant I states $B \succ A$. However, applicant I will infer the type of school A after the matching. Only if school A is of type A^1 , school A is unmatched, and the matching is immediate and ex-post unstable. \square

C.2.6 Statement 6

With reciprocal preferences, neither the TSDA nor the RPDA is strategy-proof for firms in a standard two-sided matching market with firms on the proposing side.

We show with Example C.6 that neither the TSDA nor the RPDA is strategy-proof for firms in a standard two-sided matching market if firms are on the proposing side.

Table C.6: Example 10

Proposer / Firm		Receiver / Worker
$A^1 : I \succ II$	with (p)	$I : A_1 \succ I \succ A_2$
$A^2 : II \succ I$	with $(1 - p)$	$II : B \succ II$
$B : II \succ I$		

Given:

$$I : p \cdot u(A_1) + (1 - p) \cdot u(A_2) > u(I)$$

Proof. TSDA: In equilibrium, firm A^1 states true preferences. However, firm A^2 will mimic the strategy of firm A^1 . Worker I cannot infer the type of firm A but will state firm A as her first choice due to the utility condition $[p \cdot u(A_1) + (1 - p) \cdot u(A_2) > u(I)]$. Firm B and worker II state true preferences. No player has an incentive to deviate from their strategy. Stating true preferences for player A^2 is not an equilibrium because worker I can infer the type of firm A^2 . Worker I would prefer to stay unmatched, resulting in no match for firm A^2 .

RPDA: The same logic applies to the RPDA. Assume that both receivers state their true, complete reciprocal preferences. If both types of firm A state true preferences then type A^2 is unmatched. Therefore, type A^2 will misrepre-

sent their preferences and state $I \succ II$. Firm A will always be matched with worker I .

Hence, both for the TSDA and RPDA truth-telling is no weakly dominant strategy. \square

C.2.7 Statement 7

There are markets where the DA mechanism achieves stable outcomes in undominated strategies that imply truthful reporting for firms, while the TSDA does not.

We show one exemplary matching market (Example C.7) where the DA mechanism achieves stable outcomes in undominated strategies that imply truthful reporting for firms, while the TSDA does not.

Table C.7: Example 11

Proposer / Firm		Receiver / Worker
$A^1 : I \succ II$	with (p)	$I : B \succ A$
$A^2 : II \succ I$	with (q)	$II : A \succ B$
$A^3 : I$	with $(1 - p - q)$	
$B^1 : II \succ I$	with (w)	
$B^2 : I \succ II$	with $(1 - w)$	

Given:

$$I : u(A) > \frac{p}{1-q} \cdot u(B) + \frac{1-p-q}{1-p} \cdot u(I)$$

$$II : u(B) > \frac{p}{1-q} \cdot u(A) + \frac{1-p-q}{1-p} \cdot u(II)$$

$$A^1 : u(II) < w \cdot u(II) + (1-w) \cdot u(A)$$

Proof. DA: With a DA in place, stating true preferences is an equilibrium for both sides of the market. As shown before, truth-telling is a weakly dominant strategy for proposers in a DA mechanism. Due to the assumption on the

utility of worker I and worker II truth-telling is also optimal for the workers [$I : u(A) > \frac{p}{1-q} \cdot u(B) + \frac{1-p-q}{1-p} \cdot u(I)$ and $II : u(B) > \frac{p}{1-q} \cdot u(A) + \frac{1-p-q}{1-p} \cdot u(II)$]. Receivers have no incentive to state truncated preferences because the expected utility of misrepresenting preferences is lower than the expected utility of stating true preferences. Given that all players state true preferences, the matching is stable.

TSDA: If firms state their true preferences, applicants can infer the firms' types and strategically misrepresent their preferences to ensure their preferred matching. Due to the sequential game the misrepresentation by the workers can be prevented if proposers truncate their preference list before the receivers do. Hence, there is a trade off for firms. If all firms state true preferences the receivers will misrepresent preferences in some cases and ensure their preferred matching. If firm A^1 misrepresents its preferences (truncating) it ensures a better match in some cases while it risk being unmatched in other cases. Given the assumptions on the utilities [$A^1 : u(II) < w \cdot u(II) + (1-w) \cdot u(A)$], firm A^1 truncates its preferences if being of type A^1 while firm B does not. In the illustrated market, if firm A is of type A^1 , it will truncate its list by only stating I as potential partner to prevent receivers from misrepresenting their preferences, who wish to do whenever confronted with types A^1 and B^1 . However, it involves the risk of being unmatched in the event of B^2 . In that case, the matching is immediate and ex-post unstable. \square

Bibliography

- Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, *93*(3), 729–747.
- Akerlof, G. A. (1982). Labor contracts as partial gift exchange. *Quarterly Journal of Economics*, *97*(4), 543–569.
- Akerlof, G. A., & Kranton, R. E. (2000). Economics and identity. *Quarterly Journal of Economics*, *115*(3), 715–753.
- Andreoni, J., & Vesterlund, L. (2001). Which is the fair sex? Gender differences in altruism. *Quarterly Journal of Economics*, *116*(1), 293–312.
- Antler, Y. (2015). Two-sided matching with endogenous preferences. *American Economic Journal: Microeconomics*, *7*(3), 241–58.
- Antler, Y. (2019). No one likes to be second choice. *Economic Journal*, *129*(619), 1119–1138.
- Aronson, E., & Worchel, P. (1966). Similarity versus liking as determinants of interpersonal attractiveness. *Psychonomic Science*, *5*(4), 157–158.
- Athey, S., & Levin, J. (2001). Information and competition in US forest service timber auctions. *Journal of Political Economy*, *109*(2), 375–417.
- Athey, S., Levin, J., & Seira, E. (2011). Comparing open and sealed bid auctions: Evidence from timber auctions. *Quarterly Journal of Economics*, *126*(1), 207–257.
- Ausubel, L., Burkett, J., & Filiz-Ozbay, E. (2017). An experiment on auctions with endogenous budget constraints. *Experimental Economics*, *20*(4), 973–1006.

- Avery, C., & Levin, J. (2010). Early admissions at selective colleges. *American Economic Review*, *100*(5), 2125–2156.
- Bajari, P., Houghton, S., & Tadelis, S. (2014). Bidding for incomplete contracts: An empirical analysis of adaptation costs. *American Economic Review*, *104*(4), 1288–1319.
- Barankay, I. (2012). Rank incentives: Evidence from a randomized workplace experiment. *Working Paper*.
- Basteck, C., & Mantovani, M. (2018). Cognitive ability and games of school choice. *Games and Economic Behavior*, *109*, 156–183.
- Baumeister, R. F., & Leary, M. R. (1995). The need to belong: Desire for interpersonal attachments as a fundamental human motivation. *Psychological Bulletin*, *117*(3), 497–529.
- Berlin, N., & Dargnies, M.-P. (2016). Gender differences in reactions to feedback and willingness to compete. *Journal of Economic Behavior & Organization*, *130*, 320–336.
- Bichler, M., & Paulsen, P. (2018). A principal-agent model of bidding firms in multi-unit auctions. *Games and Economic Behavior*, *111*, 20–40.
- Bikhchandani, S. (2017). Stability with one-sided incomplete information. *Journal of Economic Theory*, *168*, 372–399.
- Boss, V., Dahlander, L., Ihl, C., & Jayaraman, R. (2021). Organizing entrepreneurial teams: A field experiment on autonomy over choosing teams and ideas. *Organization Science (Forthcoming)*.
- Bradler, C., Dur, R., Neckermann, S., & Non, A. (2016). Employee recognition and performance: A field experiment. *Management Science*, *62*(11), 3085–3099.
- Burkett, J. (2015). Endogenous budget constraints in auctions. *Journal of Economic Theory*, *158*, 1–20.
- Burks, S. V., Carpenter, J. P., Goette, L., & Rustichini, A. (2009). Cognitive skills affect economic preferences, strategic behavior, and job attach-

- ment. *Proceedings of the National Academy of Sciences*, 106(19), 7745–7750.
- BVerfG, 1 BvL 3/14, Urteil des Ersten Senats vom 19. Dezember 2017 (2017). https://www.bundesverfassungsgericht.de/SharedDocs/Entscheidungen/DE/2017/12/ls20171219_1bvl000314.html
- Chakraborty, A., Citanna, A., & Ostrovsky, M. (2010). Two-sided matching with interdependent values. *Journal of Economic Theory*, 145(1), 85–105.
- Chandrasekhar, A. G., Kinnan, C., & Larreguy, H. (2018). Social networks as contract enforcement: Evidence from a lab experiment in the field. *American Economic Journal: Applied Economics*, 10(4), 43–78.
- Che, Y.-K., & Gale, I. (1998). Standard auctions with financially constrained bidders. *The Review of Economic Studies*, 65(1), 1–21.
- Chen, D., Schonger, M., & Wickens, C. (2016). Otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97.
- Chen, R., & Gong, J. (2018). Can self selection create high-performing teams? *Journal of Economic Behavior & Organization*, 148, 20–33.
- Chen, Y., & Li, S. X. (2009). Group identity and social preferences. *American Economic Review*, 99(1), 431–457.
- Chiu, Y. S., & Weng, W. (2009). Endogenous preferential treatment in centralized admissions. *RAND Journal of Economics*, 40(2), 258–282.
- Compiani, G., Haile, P., & Sant’Anna, M. (2020). Common values, unobserved heterogeneity, and endogenous entry in us offshore oil lease auctions. *Journal of Political Economy*, 128(10), 3872–3912.
- Cooper, D., & Fang, H. (2008). Understanding overbidding in second price auctions: An experimental study. *Economic Journal*, 118(532), 1572–1595.
- Currarini, S., Jackson, M. O., & Pin, P. (2009). An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, 77(4), 1003–1045.

- Currarini, S., & Mengel, F. (2016). Identity, homophily and in-group bias. *European Economic Review*, *90*, 40–55.
- Davey, C. G., Allen, N. B., Harrison, B. J., Dwyer, D. B., & Yücel, M. (2010). Being liked activates primary reward and midline self-related brain regions. *Human Brain Mapping*, *31*(4), 660–668.
- Deversi, M., Kocher, M. G., & Schwioren, C. (2020). Cooperation in a company: A large-scale experiment. *CESifo Working Paper No. 8190*.
- Dreyfuss, B., Heffetz, O., & Rabin, M. (2021). Expectations-based loss aversion may help explain seemingly dominated choices in strategy-proof mechanisms. *American Economic Journal: Microeconomics (Forthcoming)*.
- Dur, R. (2009). Gift exchange in the workplace: Money or attention? *Journal of the European Economic Association*, *7*(2-3), 550–560.
- Dur, U., Pathak, P. A., Song, F., & Sönmez, T. (2022). Deduction dilemmas: The Taiwan assignment mechanism. *American Economic Journal: Microeconomics*, *14*(1), 164–185.
- Eckel, C. C., & Grossman, P. J. (2005). Managing diversity by creating team identity. *Journal of Economic Behavior & Organization*, *58*(3), 371–392.
- Fernandez, M. (2020). Deferred acceptance and regret-free truth-telling. *Working Paper*.
- Fernandez, M., Rudov, K., & Yariv, L. (2022). Centralized matching with incomplete information. *American Economic Review: Insights*, *4*(1), 18–33.
- Fischbacher, U., Gächter, S., & Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*, *71*(3), 397–404.
- Flambard, V., & Perrigne, I. (2006). Asymmetry in procurement auctions: Evidence from snow removal contracts. *Economic Journal*, *116*(514), 1014–1036.
- Gächter, S., Johnson, E. J., & Herrmann, A. (2022). Individual-level loss aversion in riskless and risky choices. *Theory and Decision*, *92*(3), 599–624.

- Gächter, S., Kölle, F., & Quercia, S. (2017). Reciprocity and the tragedies of maintaining and providing the commons. *Nature Human Behaviour*, 1(9), 650–656.
- Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1), 9–15.
- Gerhards, L., & Kosfeld, M. (2020). I (don't) like you! But who cares? Gender differences in same-sex and mixed-sex teams. *The Economic Journal*, 130(627), 716–739.
- Gill, D., & Prowse, V. (2016). Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy*, 124(6), 1619–1676.
- Goeree, J. K., McConnell, M. A., Mitchell, T., Tromp, T., & Yariv, L. (2010). The 1/d law of giving. *American Economic Journal: Microeconomics*, 2(1), 183–203.
- Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1), 114–125.
- Grenet, J., He, Y., & Kübler, D. (2022). Preference discovery in university admissions: The case for dynamic multioffer mechanisms. *Journal of Political Economy*, 130(6), 1427–1476.
- Gundlach, J. (2021). *Per Algorithmus zum Kitaplatz? Potenziale und Erfolgsfaktoren für eine bessere Kitaplatzvergabe mithilfe von algorithmischen Systemen*. Bertelsmann Stiftung.
- Hakimov, R., & Kübler, D. (2021). Experiments on centralized school choice and college admissions: A survey. *Experimental Economics*, 24(2), 434–488.
- Hedegaard, M. S., & Tyran, J.-R. (2018). The price of prejudice. *American Economic Journal: Applied Economics*, 10(1), 40–63.

- Hendricks, K., Porter, R., & Tan, G. (1993). Optimal selling strategies for oil and gas leases with an informed buyer. *American Economic Review*, *83*(2), 234–239.
- Horwitz, S. K., & Horwitz, I. B. (2007). The effects of team diversity on team outcomes: A meta-analytic review of team demography. *Journal of Management*, *33*(6), 987–1015.
- Jain, P. (2020). Imperfect monitoring and informal insurance: The role of social ties. *Journal of Economic Behavior & Organization*, *180*, 241–256.
- Jensen, M. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review*, *76*(2), 323–329.
- Joseph, K., & Kalwani, M. (1998). The role of bonus pay in salesforce compensation plans. *Industrial Marketing Management*, *27*(2), 147–159.
- Jung, H., Kosmopoulou, G., Lamarche, C., & Sicotte, R. (2019). Strategic bidding and contract renegotiation. *International Economic Review*, *60*(2), 801–820.
- Kosmopoulou, G., & Zhou, X. (2014). Price adjustment policies in procurement contracting: An analysis of bidding behavior. *Journal of Industrial Economics*, *62*(1), 77–112.
- Kube, S., Maréchal, M. A., & Puppe, C. (2012). The currency of reciprocity: Gift exchange in the workplace. *American Economic Review*, *102*(4), 1644–1662.
- Kuziemko, I., Buell, R. W., Reich, T., & Norton, M. I. (2014). “Last-place aversion”: Evidence and redistributive implications. *Quarterly Journal of Economics*, *129*(1), 105–149.
- Lazarova, E., & Dimitrov, D. (2017). Paths to stability in two-sided matching under uncertainty. *International Journal of Game Theory*, *46*(1), 29–49.
- Lee, S., & Niederle, M. (2015). Propose with a rose? Signaling in internet dating markets. *Experimental Economics*, *18*(4), 731–755.

- Leider, S., Möbius, M. M., Rosenblat, T., & Do, Q.-A. (2009). Directed altruism and enforced reciprocity in social networks. *Quarterly Journal of Economics*, *124*(4), 1815–1851.
- Levine, D. K. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, *1*(3), 593–622.
- Li, T., & Zhang, B. (2010). Testing for affiliation in first-price auctions using entry behavior. *International Economic Review*, *51*(3), 837–850.
- Liu, Q., Mailath, G. J., Postlewaite, A., & Samuelson, L. (2014). Stable matching with incomplete information. *Econometrica*, *82*(2), 541–587.
- Malenko, A., & Tsoy, A. (2019). Selling to advised buyers. *American Economic Review*, *109*(4), 1323–1348.
- McKinney, C. N., Niederle, M., & Roth, A. E. (2005). The collapse of a medical labor clearinghouse (and why such failures are rare). *American Economic Review*, *95*(3), 878–889.
- McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, *27*(1), 415–444.
- McWhirter, R. M., & Jecker, J. D. (1967). Attitude similarity and inferred attraction. *Psychonomic Science*, *7*(6), 225–226.
- Meisner, V., & von Wangenheim, J. (2021). School choice and loss aversion. *Working Paper*.
- Möbius, M. M., Niederle, M., Niehaus, P., & Rosenblat, T. S. (2022). Managing self-confidence: Theory and experimental evidence. *Management Science* (Forthcoming).
- Montoya, R. M., & Horton, R. S. (2012). The reciprocity of liking effect. In M. Paludi (Ed.), *The Psychology of Love* (pp. 39–57). Santa Barbara, CA: Praeger.
- Montoya, R. M., & Horton, R. S. (2014). A two-dimensional model for the study of interpersonal attraction. *Personality and Social Psychology Review*, *18*(1), 59–86.

- Montoya, R. M., & Insko, C. A. (2008). Toward a more complete understanding of the reciprocity of liking effect. *European Journal of Social Psychology*, *38*(3), 477–498.
- Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, *6*(1), 58–73.
- Opitz, T., & Schwaiger, C. (2022a). Everyone likes to be liked – Experimental evidence for reciprocal preferences in matching markets. (*Unpublished*).
- Opitz, T., & Schwaiger, C. (2022b). Reciprocal preferences in matching markets. (*Unpublished*).
- Pais, J., & Pintér, Á. (2008). School choice and information: An experimental study on matching mechanisms. *Games and Economic Behavior*, *64*(1), 303–328.
- Pais, J., Pintér, Á., & Veszteg, R. F. (2011). College admissions and the role of information: An experimental study. *International Economic Review*, *52*(3), 713–737.
- Roth, A. E. (1982). The economics of matching: Stability and incentives. *Mathematics of Operations Research*, *7*(4), 617–628.
- Roth, A. E. (1984). Misrepresentation and stability in the marriage problem. *Journal of Economic Theory*, *34*(2), 383–387.
- Roth, A. E. (1989). Two-sided matching with incomplete information about others' preferences. *Games and Economic Behavior*, *1*(2), 191–209.
- Shastry, G. K., Shurchkov, O., & Xia, L. L. (2020). Luck or skill: How women and men react to noisy feedback. *Journal of Behavioral and Experimental Economics*, *88*, Article 101592.

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.

12.09.2022

Christoph Robert Schwaiger