Essays in Microeconomics

Inauguraldissertation
zur Erlangung des akademischen Grades eines
Doktors der Wirtschaftswissenschaften
der Universität Mannheim

vorgelegt von
Marcel Preuß, M.Sc.

Graduate School of Economic and Social Sciences
Center for Doctoral Studies in Economics

an der Universität Mannheim
im Frühjahr-/Sommersemester 2018

Prodekan: Professor Dr. Jochen Streb
Referent: Professor Dr. Martin Peitz
Koreferent: Professor Dr. Henrik Orzen
Tag der mündlichen Prüfung: 16.05.2018

Acknowledgments

Writing this thesis and conducting the research contained in it would have been impossible if it wasn't for the effort and support of my supervisors Henrik Orzen and Martin Peitz. I am obliged to Martin Peitz, especially for his rigorous criticism of early versions of my work and for his irreplaceable guidance in the final phase of the PhD program. To Henrik Orzen, I am most indebted for his extremely thorough remarks during our hours-long discussions that helped me structure the ideas that finally led to two chapters of this thesis and for allowing me to benefit from his great experience in experimental economics while working on these projects.

I am also very thankful to Chengsi Wang, Dietmar Fehr and Hans-Peter Grüner for their valuable and helpful comments that have also influenced the direction of this research. Furthermore, I would like to thank Andrej Svorencik and Christian König gen.Kerting for helping me organize the sessions at the experimental labs in Mannheim and Heidelberg.

Most likely, I also would have never reached this point without Lennart and Robert, two former CDSE students I shared an office and became friends with in the first year of graduate school. For working together on our problem sets and helping me handle the first year courses in macroeconomics and econometrics, I want to thank them most sincerely.

Moreover, I am incredibly grateful to Annalena, Justin, Iris, Martin and Philipp for making life outside of graduate school enjoyable and for bearing with me even in the most stressful times.

Last but not least, I wish to thank my family, and my mother Carina in particular, for their enduring support and encouragement.

Contents

List of Figures

List of Tables

1 General Introduction

As a social science, economics is concerned with the study of all kinds of social interactions. Perhaps most prominently, Garry Becker has demonstrated how applying economic analysis to a wider range of human behavior and social interactions, including nonmarket environments, constitutes a powerful tool to obtain novel insights. Following this view, the questions explored in this thesis deal with market as well as nonmarket interactions, and they are analyzed using both microeconomic theory and experimental economics as methods.

The thesis has three self-contained chapters each contributing to a different topic in microeconomics. In the second chapter, I revisit the well-known problem of price discrimination in the light of the recent rise of online tracking technologies. In an environment where consumers search for the offers of competing sellers sequentially, the ability of sellers to condition their offers on a consumer's search history has radical implications for equilibrium search behavior and prices, sometimes leading to a reversal of the theoretical predictions obtained from existing models of price discrimination. Under what conditions and why tracking of search histories is beneficial or detrimental for consumers, sellers and overall welfare are the main questions addressed in this chapter. Related to a quite different strand of research, the third and fourth chapter present and discuss findings from two experiments designed to explore the determinants of distributional preferences. While we already know a lot about what types of distributional preferences prevail in different contexts, we know relatively little about the general mechanisms underlying the formation of them. In an attempt to reduce this gap, the third chapter explores the role of people's experience of economic success on two factors possibly shaping people's preferences: their belief about the importance of luck and their fairness ideal. Another human trait underlying distributional preferences is investigated in chapter four, which deals with the potential consequences of rank-loss aversion.

In chapter 2, titled *Consumer Search with and without Tracking*, I develop a tractable framework with sequential consumer search to address the effect of tracking on market outcomes. The type of tracking this chapter focuses on is one of the most basic features of nowadays tracking technologies which is to enable websites to observe a user's browsing history through the use of third-party cookies. Tracking search histories is informative about consumers' valuations because continuing to search is costly and, thus, depends on how much a consumer likes the offers she has already encountered. Hence, given the equilibrium search rule, a consumer who has already sampled offers from multiple sellers is then more likely to have a picky taste than a consumer with an empty search history. With tracking and for a wide range of conditions, the unique equilibrium price path is increasing whereas without tracking, an average uniform price prevails. The welfare effects of this form of tracking largely depend on how tracking affects consumers' search persistence. For intermediate search costs, tracking based price discrimination exacerbates the hold-up problem and leads to inefficiently low search persistence. For high search costs instead, tracking prevents a market breakdown as low prices conditional on short search histories secure consumers a positive surplus from search. In addition, I explore whether consumers would block tracking if they could dynamically opt out from it. Interestingly, tracking prevails endogenously since disclosing their search history is always individually rational for consumers, irrespective of the overall effect on consumer surplus.

In chapter 3, titled *Heterogeneous Fairness Views*, I experimentally investigate the effect of experiencing economic success on preferences for redistribution though two channels: beliefs about the role of luck and fairness ideals. The design consists of a real effort and a spectator redistribution phase, allowing me to observe distributive preferences in the absence of any self-interest. Experiencing success (high income) raises the acceptance of inequality whereas experiencing failure (low income) reduces it both when income is assigned randomly and when it is based on relative performance. Using a treatment in which the process determining income is unobservable, I can infer beliefs about the role of luck from actual distributive choices. I find that successful spectators redistribute as if the process was random whereas unsuccessful spectators act in the opposite way. Thus, preferences are more aligned under uncertainty. The results provide strong evidence against an effective self-serving bias in beliefs about the role of luck. Instead, I find evidence in support of the claim that success makes individuals adopt a more libertarian fairness ideal.

Finally, chapter 4, titled *Relative Earnings and Rank-Loss Aversion*, presents the results of a laboratory experiment on income rank loss-aversion when a better ranking neither implies additional gains nor less disadvantageous inequality. After successful completion of a real-effort task, participants distribute additional money either to the participant ranked one position above or one position below themselves. When the difference with respect to the poorer recipient is sufficiently small, a trade-off between raising inequality and maintaining the earned income rank emerges. A theory of social preferences with income rank concerns predicts that under such circumstances, people will distribute less to the poorer recipient than they do otherwise. Moreover, participants with high degres of rank-loss aversion will distribute the majority of the additional money to the richer recipient, leading to an increase of overall inequality. The latter prediction is strongly supported in the data, which show a significant increase from 17% to 30% in the relative frequency of these inequality increasing choices. Evidence in support of a decline of the average transfer is not fully conclusive, yet the coefficients support the direction of the predictions.

2 Consumer Search with and without Tracking[1]

2.1 Introduction

In many markets, consumers learn about products and their prices only by searching different sellers sequentially. Often, the expected number of searches varies greatly across consumers as tastes and preferences are rarely homogeneous. Hence, observing a consumer's search history might convey relevant information to sellers. For instance, think of two consumers A and B (Alice and Bob), and suppose that both are looking for a suit for the job market. Since it is the first suit they buy for a long time, neither of them has a particular preference before they search and they visit the stores of different brands in a random order. While searching, Alice realizes that she is fine with almost any cut and color and thus does not need to search long. Instead, Bob finds that most cuts and colors do not suit him well, requiring him to search longer. When finally encountering his ideal suit, Bob's willingness to pay for it is, most likely, higher than Alice'. This is because Bob not only obtains utility from getting a new suit, but from having the right cut and color as well. In contrast, none of those features matter to Alice, implying that she is willing to pay less. In this environment with niche consumers like Bob and mass consumers like Alice, observing search histories may inform sellers about consumers' preferences and is thus going to provoke sellers' attention. Evidently, tracking a consumer's search process has become a widely used practice both for online and brick and mortar businesses and it will most likely be even

[1]I am very grateful to Martin Peitz, Alessandro Lizzeri, Andrew Rhodes, Nicolas Schutz, Sandro Sheliga, Thomas Tröger, Chengsi Wang and Asher Wolinsky for valuable and helpful comments.

more prevalent in the future due to exponentially improving technologies.[2] However, progress in understanding even the most general implications of tracking has been hindered by the lack of tractable models.

The questions I address in this paper are the following: What is the effect of tracking on market outcomes such as search behavior and prices? Does tracking always raise profits or can it - perhaps contrary to common wisdom - also benefit consumers while making sellers worse off? Do the welfare effects depend on the level of search costs? Finally, can tracking prevail in equilibrium if consumers possess measures to prevent it?

To address these questions, I propose a tractable framework of consumer search with tracking. Moreover, I account for consumer heterogeneity with respect to the nicheness of their taste as laid out in the introductory example. Since search with tracking compares with ordered search, the framework provides the first model of ordered search with heterogeneous consumers. Tracking search histories enables sellers to receive imperfect signals about a consumer's type and thus to learn about their preferences because stopping probabilities are type-dependent. In the baseline version of the model, there always exists a unique equilibrium with a price path that is strictly increasing in the order of search.

I evaluate the welfare consequences of tracking by comparing the tracking equilibrium with the equilibrium when tracking is not available. In general, niche types like Bob are made worse off from tracking while mass types like Alice are made better off because they are more likely to benefit from low prices at early sellers. The welfare consequences of tracking largely depend on its effect on consumers' search persistence, which is the number of sellers they are at most willing to sample if they do not encounter a sufficient match at earlier sellers. In general, tracking raises welfare if it leads to weakly higher search persistence and reduces it otherwise. For a wide range of intermediate search costs, search persistence decreases due to tracking since higher prices conditional on long search histories reduce the incentive for consumers to continue searching. However, search persistence can be lower without tracking if search costs are high. This happens when the market breaks down without tracking because the no-tracking price is inevitably too high.

[2]For example, Google places cookies on a user's computer if the retailer's website visited uses Google-Analytics for customer management, or if the retailer has joined one of the Google owned ad-networks *Doubleclick* or *Adwords*. Indeed, Mikians et al. (2012) find that both *Google-Analytics* and *Doubleclick* but also other online services providers and advertising networks such as those powered by *Facebook* or *Yahoo* are prevalent on the majority of the 200 most popular shopping websites. Mikians et al. (2012) also used automated bots to mimic different consumer types. Evidently, the bots' browsing histories had been tracked as searches for the same keyword yielded different search results and prices.

Overall, it seems without any doubt that tracking fundamentally changes market outcomes, irrespectively of the model used. Consequently, one of the most essential questions appears to be whether we should expect to encounter tracking in markets if it is not imposed exogeneously. In fact, online tracking often requires a consumer's (silent) consent. For example, a consumer must not delete her cookies to enable online retailers to observe her search history.[3] I therefore apply the novel framework of sequential consumer search to investigate whether tracking can arise endogenously. In addition to choosing a stopping strategy, consumers are able to opt out from tracking and thereby prevent sellers from observing their search history at any stage during the search process in this extension of the model. Surprisingly, the unique equilibrium outcome always exhibits full disclosure. The intuition behind this result goes back to Milgrom and Roberts (1986) and their striking unraveling argument. For any alternative equilibrium candidate in which a subset of possible search histories is not disclosed, there always exists a consumer whose search history belongs to the depicted subset and who is better off from allowing tracking.

The full disclosure prediction provides a rational explanation for why only few people delete their cookies or select the 'do not track"-request option provided by their web browsers.[4] Moreover, the analysis provides a useful benchmark for thinking about the regulation of personal data processing. Although both sellers and consumers individually prefer tracking, it may make particularly sellers worse off when it leads to a lower search persistence, which is typically the case for intermediate search costs. As even welfare might decrease due to the forgone matching surplus, there is potential for welfare-increasing regulation when the level of search costs lies in the depicted range.

The model builds on the following assumptions. Consumers must sample sellers sequentially at a cost $s > 0$ to learn about prices set by sellers and match values, which are independently distributed random variables. Consumers are ex ante heterogeneous as they draw match values from different distributions. To keep the model tractable, those distributions are simplified to two-point distributions. While one of the values is normalized to zero for all types, consumers differ with respect to their positive match value. Besides, the probability of drawing a positive match value is assumed to be a

[3]In the aftermath of several bills being introduced in the US to regulate tracking, all major web browsers integrated the option to send a "do not track"-request into their software. In addition, the European general data protection regulation law (GDPR) mandates to inform consumers when personal information is being processed. As browsing data qualifies as personal information, it requires websites to explicitly ask consumers to agree to the use of cookies. For more information about the interpretation and application of the law, also refer to the "Article 29 Data Protection Working Party" by the European Commission or the paper by Borgesius and Poort (2017).

[4]A study of German internet users from 2013 ("Maßnahmen der Internetnutzer: Digitaler Selbstschutz und Verzicht, conducted by the GfK) shows that while 70% are worried about their privacy, only 29% regularly delete their cookies.

function of the match value itself. In the main part of this paper, I assume that the matching probability is decreasing in the "conditional" (positive) match value. That is, high conditional match values coincide with low matching probabilities and vice versa. The assumption seems reasonable in markets consisting of mass and niche consumers as illustrated in the introductory example. Niche consumers have a particular taste hampering their willingness to consider most products suitable. However, once they encounter a product meeting their individual requirements, their utility from the product is relatively high. In contrast, mass consumers find most products satisfactory but only have an average willingness to pay for them.

When sellers learn about their position in a consumer's search process through tracking, search becomes perfectly ordered from their perspective. The analysis shows that consumers must then expect increasing prices in any equilibrium, leading to a simple stopping rule which lets only those consumers without previous matches continue search. Due to the interplay of consumer heterogeneity and the optimal stopping rule, expected demand from consumers with longer search histories is less elastic and, thus, prices indeed increase in the order of search. Consumer heterogeneity also leads to novel predictions regarding the effect of search costs on prices. As intuition suggests, consumers' search persistence is weakly decreasing in the level of search costs. However, the fact that consumers sample fewer sellers does not imply reduced competition and higher prices. From a seller's perspective, the probability of facing a consumer with a longer search history decreases while the probability of facing a consumer with a short search history increases when consumers sample fewer sellers. As demand from the latter group is more elastic, the equilibrium price is decreasing in search costs. This counterintuitive result provides a theoretical explanation for the empirical finding that low search cost environments like the internet sometimes lead to higher prices (Ellison and Ellison, 2014).

The dynamics described above help to understand the welfare implications of tracking. Generally, reduced asymmetric information due to tracking has diametrically opposed effects on consumer surplus. On the one hand, lower prices for short search histories have a positive market expansion effect. On the other hand, improved price discrimination reduces surplus from niche consumers who search longer. Importantly, tracking is not detrimental to every niche consumer per se due to the *order effect*. In other words, both a mass type like Alice and a niche type like Bob can benefit from lower prices at the beginning of search under tracking. Hence, the detrimental effect of price discrimination on consumer surplus is mitigated because learning about a consumer's types does not take place instantaneously but sequentially. If search costs are negligible such that consumers without a match never stop their search before they have sampled all sellers, tracking can be favorable both for consumers and sellers.

Otherwise, tracking affects consumers' search persistence and changes consumers' and sellers' surplus in diametrically opposed ways. The longer a consumer's search history, the larger the share of the matching surplus sellers can extract with tracking. Therefore, the expected surplus from sampling an additional seller necessarily falls below search costs beyond some fixed search history. Because the hold-up problem prevents an equilibrium with lower prices, search persistence decreases due to tracking. This effect is most pronounced for an intermediate level of search costs and implies less matches from consumers like Bob with a high willingness-to-pay, leaving sellers with less profits and reducing total welfare while the average consumer is still better off from tracking due to lower initial prices.

In contrast, tracking may also raise everyone's surplus for high search costs. Unless consumers sample only a single seller under search without tracking, the no-tracking price always exceeds the price set by the first seller under search with tracking. Yet, sampling only a single seller in equilibrium may not be consistent with the stopping rule for any level of search costs. Then, there is a range of high search costs where the market with tracking is still active whereas it breaks down without tracking, thus implying significant welfare losses from unrealized matches. The reason why a no-tracking equilibrium with a search persistence of one or few sellers may not exist is because of the adverse effects of search persistence on prices. When the expected surplus from search is negative given the level of search costs and the no-tracking price, a lower search persistence could reduce the price and make sampling the first seller worthwhile again. However, stopping search as early as implied by a low search persistence is not sequentially rational given the reduced price. Moreover, randomizing over sampling the first seller cannot change sellers' belief about the average search history in the market, and, thus, cannot restore the equilibrium with active search if this inconsistency problem prevails at the beginning of the search process.

In an extension, I study the complementary case of a matching probability function which is weakly increasing in the conditional match value. Such a positive relationship is likely to prevail in markets where consumer heterogeneity is mainly determined by heterogeneous budget sets rather than differences in taste. Hence, the two distinct cases of an in- and decreasing matching probability function refer to markets that are likely to be inherently different from one another. Also, market outcomes stand in stark contrast to the baseline version of the model since the price is now weakly decreasing in the order of search. The comparative statics of tracking compare with those for a decreasing matching probability, except for the market breakdown result, which cannot obtain in this specification. Besides, the result that tracking arises endogenously is robust to this extension of the model.

2.2 Literature

The paper relates to two broad strands in the literature. On the one hand, the search framework developed contributes to the literature on consumer search by embedding ex ante consumer heterogeneity into a model of ordered search, two areas, which so far have been studied only in separation. On the other hand, this paper studies a consumer's privacy data disclosure problem in a sequential search framework and thus relates to the literature on the economics of privacy. More precisely, it complements existing research on the consumer's data protection problem which has mostly been addressed using static models with exogenous data as opposed to a dynamic search environment with endogenous data contained in search histories. I first review the literature on consumer search before providing an overview about the paper's relation to the economics of privacy literature.

In the seminal paper by Diamond (1971), consumers search for prices in a random order. Despite multiple sellers producing a homogeneous good, sellers charge the monopoly price because demand is completely inelastic for any price below the monopoly price due to the hold-up problem. As a consequence, consumers, rationally expecting monopoly prices, are better off from not searching at all. Wolinsky (1986) shows that this counter-intuitive result, often referred to as the *Diamond Paradox*, disappears when products are differentiated and consumers thus search not only for prices but product fit as well. Anderson and Renault (1999) complement Wolinsky (1986) by showing how both the *Diamond Paradox* and the Bertrand outcome arise in the limit as either the degree of product differentiation or the level of search costs vanishes. The model by Wolinsky (1986) with the extension of Anderson and Renault (1999) (henceforth WAR) has since then become the workhorse model of consumer search for many researchers. As in their model, I assume that consumers' preferences are heterogeneous by modeling match values as independently and identically distributed shocks. As opposed to WAR however, I assume that these shocks are identically distributed only for a particular consumer type but differently distributed across types. That is, consumer heterogeneity is revealed not only ex post after sampling sellers, but already prevails *ex ante* even before search begins, thus leading to type-dependent search behavior.

Even though consumers sample sellers in no particular order in my model, tracking enables sellers to learn about their position in a consumer's search proces, which is the defining assumption in the research on ordered search and thus perhaps the closest literature this paper relates to. Arbatskaya (2007) shows that search cost heterogeneity can explain why prices might increase in the order of search even when consumers search for a homogeneous product. Zhou (2011) considers ordered search in the WAR model with differentiated products. He also finds that prices increase in the order of search in equilibrium for the reason that later sellers possess a larger monopoly power over remaining consumers. As prices depend on a seller's position, which in turn, can be inferred perfectly from a consumer's search history, Zhou (2011), in fact, also studies search history-based price discrimination. In a related work, Armstrong and Zhou (2010) analyze a seller's optimal strategy to discriminate between "fresh" and returning consumers. In contrast to this paper, which allows for search histories of arbitrary length, they restrict attention to a duopoly version of the WAR model and mainly focus on a seller's incentive to deter consumer search by offering buy-now discounts.

Besides, even though Zhou (2011) provides a solution to the WAR model with ordered search for a specific distribution of match values, the model is not tractable enough to account for additional consumer heterogeneity that might create potentially countervailing effects regarding price discrimination and search behavior. Indeed, I find that tracking often leads to more efficient search by raising consumers' search persistence, which stands in stark contrast to Zhou's finding that ordered search leads to inefficiently low search. Moreover, his very intuitive price dispersion result disappears if the number of sellers grows and the difference in monopoly power becomes arbitrarily small due to the infinite number of remaining sellers. That is, ordered search or tracking plays no role when the number of sellers is large, suggesting that the WAR model might not capture all important aspects of search markets.

Several papers building on the WAR model and focusing on particular applications of ordered search share this property with the work by Zhou (2011). Among others, Armstrong et al. (2009), Haan and Moraga-González (2011) and Moraga-González and Petrikaitė (2013) have studied how higher quality, more advertising, or merging with a competitor can make a subset of sellers more salient and thereby lead to partially ordered search. Importantly, search in these models is ordered only with respect to the first, salient seller as it would otherwise become intractable.

Though quite different from this paper, some authors have also explored in detail applications of ordered search that are more closely related to internet search and tracking. In Chen and He (2011), a monopoly platform uses an auction to determine the order in which sellers are shown to consumers. As they assume that some sellers' products are more relevant to consumers than others, the mechanism places those sellers at the top consumers have the highest valuations for. De Cornière (2016) studies a platform using a keyword matching mechanism to determine a consumer's search order. If sellers choose to be associated with a keyword the consumer entered and pay an advertsising fee, they obtain a prominent position in the search list. In both papers, the authors start from the fact that sellers vary in terms of their relevance to particular consumers. While it is modeled explicitly only in De Cornière (2016), they thus presuppose that consumers are somewhat ex ante heterogeneous. However, tracking occurs at a single instant prior to actual search in their papers whereas it is modeled as a dynamic process enabling sellers to learn gradually about consumers from search histories in my framework.

Other papers study ex ante consumer heterogeneity within the the WAR model more explicitly but (have to) restrict attention to random search. Moraga-González et al. (2017) study price formation when consumers have different search costs. As changes in the distribution of search costs affect both the extensive and intensive search margin, they find that lower search costs can increase the price charged from actively searching consumers.[5] To understand the effect of targetability or in other words, the quality of search, Yang (2013) studies a seller's choice whether to serve a mass product many consumers like or a niche product appealing only to few. By assuming that consumers draw positive match values only from sellers serving their preferred category and that the probability of encountering a respective seller depends both on the quality of search and the category's coverage in the market, he shows how the long-tail effect is driven by the quality of search. Bar-Isaac et al. (2012) also study the long-tail effect, but do not model ex ante heterogeneity on the side of consumers. Though conceptually similar, the approach taken by Yang (2013) to model mass versus niche consumers is different from mine. Instead of introducing different product categories for mass and niche consumers, I assume that mass consumers are more likely to find any product suitable due to their less restrictive taste compared to niche consumers.

While *search history* based pricing has received relatively little attention in the literature, *purchase history* based price discrimination has been studied extensively. In the absence of online shopping and related privacy concerns, this literature with early works by Hart and Tirole (1988), Fudenberg and Tirole (2000) and Villas-Boas (1999) deals with consumers making purchase decisions in multiple periods. As a consumer's revealed choice for a particular seller is a signal of her willingness to pay, it affects prices she obtains in future periods. A common prediction in this literature is *consumer poaching*: a seller's strategy to offer low prices to those consumers who have revealed their preference for the competitor's product in the previous period. As a consequence of low prices in later periods, prices also become competitive in the initial period. However, when accounting for the possibilty of strategic waiting, Chen and Zhang (2009) identify a novel and opposing incentive to price high in the initial period

[5] I discuss the effect of search cost heterogeity in my framework in section 7.

as it allows to better learn about the willingness to pay of those consumers who still make a purchase.

More closely related to the economics privacy, Taylor (2004) considers purchase history-based customer lists as valuable information one firm would want to sell to another firm if a consumer's valuations for different products are correlated. He finds that privacy protection policies are necessary if consumers are naive, but not so otherwise as the willingness to pay of sophisticated consumers in the first period decreases when anticipating exploitation in future periods. Acquisti and Varian (2005) reaches similar conclusions in a monopoly model where they also model the consumer's decision to remain anonymous as in my model. Conitzer et al. (2012) also study the consumer's privacy choice in a model similar to Acquisti and Varian (2005) but introduce a cost to maintain anonymity.

Motivated by the rising concern for internet privacy, a number of authors have revisited the consumer's data protection problem and extended the analysis of Conitzer et al. (2012). Taylor and Wagman (2014) compare the welfare implications of maintaining privacy across different oligopoly models and find ambiguous effects. In Montes et al. (2017), competing firms can buy data containing consumers' private information from an intermediary unless consumers pay a "privacy cost" to remain anonymous. Since paying the cost also reveals some private information, a higher "privacy cost" may in fact raise consumer surplus. Belleflamme and Vergote (2016) study an environment where a monopolist is able to detect private information with some probability unless consumers use a costly technology to maintain privacy. Similarly to the previous authors, they find that the availability of such a technology makes consumers worse off. The results I obtain are quite different. Although more privacy protection would yield a higher surplus in some cases, consumers never use the technology despite its availability at no cost. Besides, the papers mentioned above apply a static model where information about consumers exists from an exogenous source. If sellers can rely on the availability of informative big data, that approach may be quite accurate. However, if informative data is rare, sellers might pay more attention to a consumer's search history for the product they sell, as studied in this paper.

2.3 Search Model

There is a continuum of consumers $i \in [0, 1]$ and a finite number of firms N. Firms are selling a horizontally differentiated product. The goods can be produced at constant marginal cost, which is normalized to zero. Sellers set their prices and can condition them on different search histories, if these are observed.

Consumers. Consumers sample sellers sequentially in no particular order and with free recall. They search for both prices and product fitness and pay a sampling cost $s > 0$ for each seller. Consumer i obtains utility $u_{ik} = v_{ik} - p_k$ if she purchases from seller $k \in \mathbb{N}$, where p_k is the price and v_{ik} captures her seller specific match utility. Match utilities for sellers' products $(v_{i1}, v_{i2}, ..., v_{iN})$ are random draws from the set $v_{ik} \in \{0, x_i\}$ and independently and identically distributed across sellers. The conditional match value x_i defines consumer i's type. A type is randomly drawn ex ante from the compact set X with $\underline{v} = \inf(X)$ and $\bar{v} = \sup(X)$ from the log-concave distribution $F(x)$. To avoid corner solutions, assume that $\underline{v} > 0$ is sufficiently small. Denote by $g(x) = Prob(v = x)$ the matching probability, which is type x_i's probability of drawing a positive match value $v_{ik} = x_i > 0$ at a random seller k. The matching probability function $g(x)$ is asumed to be log-concave and monotone decreasing, implying that high conditional match values correspond to low matching probabilities and vice versa. This assumptions seems to fit well many markets where the main difference between consumers is the extent to which they care about all features of a product. In the introductory example, Bob cares both about the cut and the color of his new suit while Alice does not. Consequently, Bob has a lower matching probability than Alice. However, his conditional match value is higher because he obtains utility from all features of the suit. A decreasing matching probability function extends this heterogeneity between picky niche consumers and accepting mass consumers to a continuous type space. Section 2.5 contains an extension of the baseline model to the case of a weakly increasing matching probability function.

Additionally, both consumers and sellers know only the distribution of types $F(\cdot)$ and a particular consumer's type x_i must be learned during search. Notice that this assumption holds in the example concerning Alice and Bob looking for a new suit. Neither of them knows about their preferences over suits ex ante. Instead, they find out about how picky they are while searching. While this simplifying assumption may seem stark, it actually renders the purpose of search more realistic. In other words, because consumers do not know their conditional match value, they truly search for both price and product fit as in the WAR model. Technically, consumers must not have perfect information about their type x_i ex ante to prevent a market breakdown result as in the *Diamond Paradox.* [6] If all consumers knew their types, sellers would always have an incentive to deviate from any price leaving a strictly positive surplus to actively searching consumers. This is because sellers know that any actively searching consumer's conditional match value must exceed the expected price as the consumer

[6] If information is sufficiently imperfect, the hold-up problem has no bite and consumers will find search worthwhile. Thus, no information about x_i is a simplifying but not strictly necessary assumption.

would not have incurred the search costs otherwise. Hence, in the only equilibrium with perfectly informed consumers, no consumer would search.

While the random match value framework I use is more stylized than in WAR, it allows me to handle the complexity arising from incorporating consumer heterogeneity into a model of search with tracking. Recently, several authors have passed on the continuous match value distribution (for individual consumers) as well in order to gain tractability, see for example Chen and He (2011), Anderson and Renault (2015) or Armstrong and Zhou (2011).

Search history. The search history h is what other sellers observe from an arriving consumer under search with tracking. I assume that other sellers cannot observe the price of previously sampled products. Further, nothing can be learned from knowing a particular sellers' identity since match values are uncorrelated across sellers. Consequently, the total number of past sellers a consumer has sampled is a sufficient statistic to update the posterior beliefs about her type. Hence, $h \in \mathbb{N}$.

Timing. Players move in the following order. First, sellers set prices conditional on any feasible search history. Under search without tracking, they set an unconditional price. Prior to searching, nature draws each consumer's type. Next, consumers search by sampling sellers sequentially at a cost s per seller. Under search with tracking, sellers observe the consumer's search history when being sampled. Consumers observe the respective price and match value and decide whether to purchase, to return to a previous seller, or to stop search.

Equilibrium concept. The equilibrium notion I consider is perfect Bayesian Nash Equilibrium (PBE). Sellers choose pricing strategies to maximize expected profits given other sellers' prices and consumers' stopping rule. Consumers maximize surplus by choosing an optimal (possibly non-stationary) stopping rule. Both sellers' beliefs about a consumer's type and consumers' expectations regarding prices need to be consistent with the equilibrium stopping rule and equilibrium pricing. Since equilibrium pricing strategies will be deterministic, I assume that consumers have passive beliefs if they observe a non-equilibrium price.[7] I restrict attention to symmetric equilibria.

[7]This restriction seems reasonable as an individual consumer is of mass zero here. Hence, an off-equilibrium discriminatory price for a single consumer followed by an out of equilibrium action by this very consumer does not change expected demand at any other seller and thus gives no rise to expect subsequent prices to be different from the equilibrium prices.

2.3.1 Search With Tracking

I begin by analyzing the equilibrium under search with tracking and search history based price discrimination and then turn to the analysis of search without tracking.

When consumers sample sellers sequentially, their decision about when to stop searching not only depends on the available match values and prices, but also on prices they expect at forthcoming sellers. The main result of this section is that that there is a unique PBE in which prices satisfy $p_1 < p_2 < ... < p_N$. Therefore, I seek to construct such an equilibrium first and assume that consumers expect $p_1^e \leq p_2^e \leq ...p_N^e$. Second, I show that these beliefs are the only beliefs permissible under rational expectations.

In the following analysis, it facilitates notation to write "seller k" when referring to the k's seller a consumer has sampled. The index k thus does not denote a specific seller for all consumers. Moreover, note that prior to sampling seller k, the consumer's history h equals $k-1$ while it is $h = k$ thereafter. Besides, I omit the subscript i for brevity when it does not lead to ambiguous statements.

Optimal stopping. If $v_{ik} = x_i$ at some seller k, a consumer has no incentive to continue to search as she expects at most to obtain x_i again but to pay a higher price at any forthcoming seller. Therefore, consumers encountering a match $v_{ik} = x_i$ buy if $x_i \geq p_k$ and stop searching without making a purchase otherwise. Whether consumers prefer continuing to search after a history of h unsuccessful matches depends on the continuation value from sampling seller $h+1$, denoted by V_{h+1}. Then, a consumer with history h samples seller $h+1$ if both $v_h = 0$ and $V_{h+1} > 0$.[8]

Lemma 1 *After inspecting seller k, the following non-stationary stopping rule, denoted by $\mathcal{R}^*$, is optimal: Buy if $v_k = x \geq p_k$ and continue to search if $v_k = 0$ and $V_{k+1} > 0$. Otherwise, end search.*

Learning. A consumer perfectly learns her type only upon encountering a match but can learn from the length of her search history otherwise. The optimal stopping rule $\mathcal{R}^*$ implies that a consumer who has a history h and who is about to sample seller $h+1$ must have received only $v_{ik} = 0$ at any seller k, $k \leq h$. As the probability of not encountering a match at h previous sellers varies with x, the search history h is informative about one's type. The expected probability of no match at a single seller is given by:

$$\int_{\underline{v}}^{\bar{v}} \big(1 - g(t)\big) f_k(t)\mathrm{d}t,$$

[8]Since V_{h+1} has no effect on equilibrium prices, its derivation is postponed to the end of this section. For now, it is sufficient to note that V_{h+1} depends on h but not on an individual's type x_i because consumers do not know their type perfectly but must learn about it from searching.

where f_k is seller k's posterior belief about a consumer's type conditional on $h = k - 1$. Since consumers observe their own history by construction, f_k represents their belief prior to sampling seller k as well. Note that in the following analysis, the initial prior $f(x)$ without subscript refers to the distribution of types expected by the first seller a consumer samples. By repeated use of Bayes' rule, I obtain the posterior belief $f_k()$ for any seller k:

$$f_k(x) = \frac{\big(1 - g(x)\big)^{k-1} f(x)}{\int_{\underline{v}}^{\bar{v}} \big(1 - g(t)\big)^{k-1} f(t)\mathrm{d}t}. \tag{2.1}$$

Consumer demand. The consumer's optimal stopping rule implies that conditional on a match, a consumer always buys the product from seller k immediately if $x_i \geq p_k$ and $p_k \leq p_j^e \; \forall \; j > k$. If the latter constraint is not binding in equilibrium, expected demand from a consumer who is known to be visiting her first firm writes:

$$D_1(p) = P(v \geq p) = \int_p^{\bar{v}} g(x) f(x)\mathrm{d}x,$$

Based on the posterior $f_k()$, a general expression of seller k's demand, denoted by $D_k(p)$, obtains:

$$D_k(p) = \int_p^{\bar{v}} \frac{\big(1 - g(x)\big)^{k-1} g(x)}{\int_{\underline{v}}^{\bar{v}} \big(1 - g(t)\big)^{k-1} f(t)\mathrm{d}t} f(x)\mathrm{d}x \tag{2.2}$$

Pricing. Seller k's demand at price p is given by (2.2) if $p \leq p_j^e \; \forall \; j > k$ since consumers might follow an alternative stopping rule otherwise. The following analysis will show that the constraint is not binding at the equilibrium price and that deviating to a price $p_k > p_j^e$ cannot be profitable neither. That is, the profit-maximizing price is independent of all competitors' prices and seller k's problem is equivalent to a monopolist's pricing decision:

$$p_k \in \arg\max_p D_k(p)p \tag{2.3}$$

The reason why monopoly prices prevail in the presence of competing sellers is similar to Diamond (1971), even though he considers consumer search for homogeneous products. Despite product differentiation however, expectations about future prices suppress any form of price competition between sellers. This holds for any strictly positive search friction. The solution to (2.3) yields:

Lemma 2 *The profit-maximizing price is uniquely defined for every seller k. The sequence of profit-maximizing prices $\{p_k^+\}_{k=1,\ldots,N}$ satisfies $p_1^+ < p_2^+ < \ldots. < p_N^+$.*

All omitted proofs are presented in appendix A. Intuitively, prices increase because consumers with a higher conditional match values need to sample more sellers on average than consumers with a low conditional match value until they encounter the first match. Therefore, the relative share of consumers with a high conditional match value is larger for longer search histories. Consequently, expected demand becomes more inelastic and profit-maximizing prices increase in a consumer's search history. The uniqueness of prices is due to the fact that the RHS of the FOC $p = -\frac{D(p)}{D'(p)}$ is decreasing in p, which follows from the log-concavity assumption about $f_k(x)$ and $g(x)$.[9]

The sequence of increasing prices $\{p_k^+\}_{k=1,\ldots,N}$ obtained in lemma 2 is optimal conditional on consumers expecting an increasing price path. Consequently, consumers follow the stopping rule $\mathcal{R}^*$ by lemma 1. Since lemma 2 shows that $\{p_k^+\}_{k=1,\ldots,N}$ is the profit-maximizing sequence of prices given $\mathcal{R}^*$, an equilibrium with prices $\{p_k^+\}_{k=1,\ldots,N}$ and consumer stopping characterized by $\mathcal{R}^*$ indeed exists. Nevertheless, there might be other equilibria. Instead of expecting an increasing price path, consumers might initially expect a decreasing or non-monotonic price path, leading to a different stopping rule and thus to different prices. However, even when allowing for arbitrary consumer beliefs about prices along their search path, the only beliefs consistent with equilibrium pricing of sellers are those of an increasing price sequence given by $\{p_k^+\}_{k=1,\ldots,N}$.[10] This can be shown by means of contradiction. Suppose that consumer expectations $\{p_k^e\}_{k=1,2,\ldots,N}$ are not increasing in k. First, note that:

Lemma 3 *In any PBE, consumer expectations satisfy $p_k^e \geq p_k^+ \ \forall \ k \leq K^*$.*

To see intuitively why lemma 3 holds, consider seller j^* where j^* denotes the seller closest to the end of the search process whose expected price $p_{j^*}^e$ lies below $p_{j^*}^+$. The fact that all remaining sellers are expected to charge higher prices by construction has important implications for seller j^* when deviating to a price in the neighborhood of $p_{j^*}^e$. Seller j^*'s expected demand does not depend on whether arriving consumers have available matches from previous sellers since sampling j^* is only worthwhile if they are in fact willing to buy at $p_{j^*}^e$. That is, seller j^* can sell to all consumers whose

[9] Bagnoli and Bergstrom (2005) discuss properties of log-concave functions and show that the FOC for demand functions of the type $D(p) = \int_p^{\bar{v}} h(x)\mathrm{d}x$ is decreasing in p if $h(x)$ is log-concave. Hence, decreasingness follows from log-concavity of both $g(x)$ and $f_k(x) \ \forall \ k$.

[10] As in other models of consumer search, there always exists an uninteresting equilibrium where consumers expect prices to be larger than their expected surplus from search. In such an equilibrium, no consumer searches and setting such high prices indeed constitutes an equilibrium strategy for sellers.

match value exceeds the price and thus has full monopoly power over its demand, implying that his problem can be characterized by (2.3). Nevertheless, the profit maximizing price need not be equal to p^+_{j*} due to changes in the distribution of arriving consumers. This is because consumers may have applied a stopping rule different from $\mathcal{R}^*$at previous sellers.

Changes in the distribution of arriving consumers must be of the following kind: If a consumer samples j^* despite an available match, her conditional match value must satisfy $x_i > p^e_{j*}$ as she already knows her conditional match value and would otherwise be better off from not sampling j^*. Hence, if due to any alternative stopping rule demand at seller j^* changes, it is due to an increase in demand from types $x_i > p^e_{j*}$. Notably, demand of these additional consumers attracted by the lower expected price is completely inelastic in the neighborhood of p^e_{j*}. Moreover, note that lemma 2 implies that any (local) upward deviation to $p'_{j*} > p^e_{j*}$ is profitable for seller j^* even in the absence of those additional consumers. Hence, setting a price above p^e_{j*} is for sure profitable in the presence of this additional, perfectly inelastic demand. Since this is true for any $p^e_{j*} < p^+_{j*}$, the argument can be applied repeatedly from the last to the first seller, yielding lemma 3.

In addition, notice that consumers apply $\mathcal{R}^*$ at the first seller if he sets a price p_1 in the neighborhood of p^+_1 since $p^e_k > p^+_1$ by lemma 3 $\forall\ k > 1$. By lemma 2, p^+_1 maximizes the first seller's profits under the stopping rule $\mathcal{R}^*$. Hence, any alternative price $p'_1 > p^+_1 + \delta$ ($\delta > 0$) inducing an alternative stopping rule for some types must yield strictly lower profits. This is because under any alternative stopping rule, there are some types who continue searching despite a match $v_{i1} > p_1$ and return only with some probability less than one. Consequently, demand and thus profits must be strictly lower than when consumers apply $\mathcal{R}^*$. It follows that consumers must expect $p_1 = p^+_1$ in any PBE. Further, the same argument applies to the second seller a consumer visits and so forth. Thus, no equilibrium exists, in which consumers do not expect an increasing price path.

Proposition 2.1 *With tracking, in the unique equilibrium, prices increase in the order of search and consumers follow the stopping-rule $\mathcal{R}^*$. The equilibrium always exists.*

Search persistence. While search costs have no effect on equilibrium prices under search with tracking, they matter for consumers' search persistence: how long to continue search if no match occurs. Consumers' search persistence is captured by $K^* \in \mathbb{N}$. First, note that K^* depends on a consumer's continuation value V_{h+1} at any history h. Its recursive formulation writes:

$$V_{h+1} = \mathbb{E}\big[\max\big\{[v_{h+1} - p_{h+1}], 0, V_{h+2}\big\}\big] - s.$$

To derive an explicit expression for V_{h+1}, one needs to account for K^* since the continuation value from sampling seller k must contain the option value from continuing to search at least seller $k+1$, which is feasible only if seller k is sampled first. For any K^* and history $h < K^*$, the continuation value writes:

$$V_{h+1}(K^*, s, p_1, p_2, ..., p_{K^*}) = \frac{1}{\int_{\underline{v}}^{\bar{v}} \big(1-g(x)\big)^h f(x)\mathrm{d}x} \cdot \Big\{ ...$$

$$\sum_{j=h+1}^{K^*} \Bigg(\underbrace{\int_{p_j}^{\bar{v}} g(x)(x-p_j)\big(1-g(x)\big)^{j-1} f(x)\mathrm{d}x}_{\textit{weighted matching surplus}} - s \underbrace{\int_{\underline{v}}^{\bar{v}} \big(1-g(x)\big)^{j-1} f(x)\mathrm{d}x}_{\textit{expected search attmepts}} \Bigg) \Bigg\}. \quad (2.4)$$

Even though $V_{h+1}(K^*, s, p_1, ..)$ always depends on all sellers' prices, equilibrium search persistence K^* and search costs s, I omit those arguments for brevity when it does not affect comprehensibility. The term before the curly brackets is part of the belief updating regarding the consumer's type and for $h = 0$, the term disappears. It equals the inverse of the total probability of not finding a match after sampling h sellers and thus normalizes the probability of encountering a match at sellers $h+1$, $h+2$, and so forth. The term within the brackets sums over the surpluses from additional search attempts weighted by the updated consumer's type after history $h+j$. Notice that the sum of additional search attempts goes from any history h to K^*.

For the optimal search persistence K^*, it must hold that $V_k(K^*) \geq 0 \; \forall \; k \leq K^*$ and $V_{K^*+1}(K^*+1) < 0$. In words, sampling any seller prior to K^* must be rational. Moreover, there cannot be another $K' > K^*$ satisfying the first the first condition and rendering $V_{K^*+1}(K') \geq 0$. Since continuation values are affected by search costs, K^* depends on search costs as well. Formally, define $\hat{K}(s) := \{K : V_k(K, s) \geq 0 \; \forall \; k \leq K \in \mathbb{N}\}$. Then, equilibrium search persistence satisfies $K^* \in \mathbb{K}(s)$, where

$$\mathbb{K}(s) := \big\{K : V_{K+1}(K', s) < 0 \; \forall \; K' \in \big(\hat{K}(s) \cap \{k : K' > K\}\big)\big\}. \quad (2.5)$$

By construction of $\mathbb{K}(s)$, there always exists exactly one K^* for any given s. Besides, equation (2.4) shows that for a fixed K^*, $V_{h+1}(K^*, s)$ is decreasing in s for any history h. Hence, $\hat{K}(s') \subseteq \hat{K}(s)$ for any $s' > s$. Consequently, $K^*(s)$ as defined in equation (2.5) must be weakly decreasing in s. This result is, of course, very intuitive. As prices are independent of search costs, an increase in search costs reduces the continuation value for any history h. That is, an increase in search costs can only make consumers switch from continuing to search given a particular history h to stopping to search given the same history, but will never induce a change in the other direction. As a consequence, search persistence cannot be increasing in search costs.

2.3.2 Search Without Tracking

Without tracking search histories, sellers cannot price discriminate. Therefore, sellers expect the same demand (elasticity) from any newly arriving consumer. Since given those expectations, only one price maximizes profits, search without tracking implies a uniform equilibrium price set by all sellers. As a consequence, consumers have no incentive to defer the purchase decision after encountering a match and thus their optimal stopping rule equals $\mathcal{R}^*$. That is, i either buys if $x_i > p_{ij}$ or leaves the market without a purchase which is identical to the uniquely optimal stopping rule when $g(x)$ is decreasing. Consequently, sellers set prices monopolistically. Due to the simple optimal stopping rule, the distribution of consumer types $f_k(x)$ at any seller k is correctly specified by the posterior belief derived in (2.1). Further, demand functions - if sellers could discriminate - are given by (2.2), i.e. they are identical to those under search with tracking. However, only consumers but not sellers can update the beliefs conditional on different search histories. Hence, the sellers' expected demand is composed of the expected demand for each possible search history, weighted by the respective probabilities. Based on the common prior $F(x)$ and the matching probability function $g(x)$, sellers can compute the probability that a consumer has a history of $h = 0, 1, 2, ..., K^* - 1$ previous sellers. Note that consumers' search persistence K^* depends only on equilibrium prices but does not change in response to any deviating price. I analyze the equilibrium level of K^* after discussing the equilibrium pricing. The probability ϕ_k of being in position $k = h + 1$ in a consumer's search process writes:

$$\phi_k = \tag{2.6}$$

$$\underbrace{\frac{1}{N} \cdot \prod_{j=1}^{k-1} \int_{\underline{v}}^{\bar{v}} \big(1 - g(t)\big) f_j(t) \mathrm{d}t}_{\textit{no match up to } k-1} = \frac{1}{N} \cdot \prod_{j=1}^{k-1} \int_{\underline{v}}^{1} \frac{\big(1 - g(t)\big)^{j} f(t)}{\int_{\underline{v}}^{\bar{v}} \big(1 - g(t)\big)^{j-1} f(t) \mathrm{d}t} \mathrm{d}t$$
$$= \frac{1}{N} \int_{\underline{v}}^{1} \big(1 - g(t)\big)^{k-1} f(t) \mathrm{d}t \;\forall\; k \leq K^*.$$

and $\phi_k = 0 \;\forall\; k > K^*$.

Notice that ϕ_k is the unconditional probability for being in a particular position. While this is the actual probability of being sampled by a consumer with history $h = k - 1$, sellers can condition the probability on the fact that the consumer is still searching. However, normalizing by $1/\sum_k^{K^*} \phi_k$ has no effect on a seller's first order condition. Expected demand is composed of the expected demand functions for each possible search history, weighted by the respective probabilities. Denote by $D(p)$ the expected demand weighing the individual demand functions from $D_1(p)$ to $D_{K^*}(p)$ at the seller's non-discriminatory unit price p, i.e. $D(p) = \sum_{k=1}^{K^*} \phi_k D_k(p)$. Sellers maximize:

$$\Pi(p) = p \sum_{i=1}^{K^*} \phi_i D_i(p) = \sum_{i=1}^{K^*} \phi_i D_i(p) p = \sum_{i=1}^{K^*} \phi_i \pi_k(p) \tag{2.7}$$

where $\pi_k(p)$ equals the profit function of a seller at the k's position if discrimination was feasible and where ϕ_i is given by (2.6) (see the appendix) and $D_k(p)$ by (2.2). As log-concavity of the individual demand functions is preserved in this weighted demand function, I obtain the following:

Proposition 2.2 *Without tracking, the unique equilibrium has a uniform price.*

Since sellers cannot observe consumers' strategies before consumers make a purchase, they must have symmetric beliefs in equilibrium. Given these beliefs, there exists a unique optimal price, set by all sellers. Thus, consumers must believe that prices are constant in any PBE. Consequently, the equilibrium with the price maximizing (2.7) is unique even when allowing for arbitrary expectations ex ante.

Without tracking, there is no price discrimination. Yet, consumer heterogeneity has another effect on the comparative statics of equilibrium price, which depends on consumers' search persistence. Intuitively, the more sellers consumers are at most willing to sample, the smaller the share of consumers with short search histories each seller can expect because probability mass shifts to longer search histories. Since demand from consumers with longer search histories is less price elastic, the equilibrium price increases. Let $p(K^*)$ be the uniform random search price if a consumer's search persistence equals K^*, then:

Lemma 4 *For $K_2^* > K_1^*$, it holds that $p(K_2^*) > p(K_1^*)$.*

By lemma 4, more persistent search behavior by consumers leads to higher prices. Thus, consumer heterogeneity implies a novel, and perhaps surprising, effect of search persistence on the equilibrium price, which is not present in the WAR model.

2.3.3 Search Persistence under Search Without Tracking

An increase in K^* as discussed in lemma 4 can be the result either of an increase in the number of sellers in the market or of a reduction in search costs. The former is true if search costs are sufficiently low such that $K^* = N$. Intuitively, the latter might be true if $0 < K^* < N$ and K^* increases due to a decrease in search costs. While this turns out to be correct, it does not follow immediately from the construction of $K^*(s)$ given in equation (2.5) due to the reverse effect of K^* on prices.

The full characterization of equilibrium search persistence as a function of search costs is provided by the lemmata 13, 14 and 15 in appendix A. To summarize, there exist two disjoint sets of search cost intervals that give rise to a different characterization of search persistence. By lemma 13, consumers' search persistence K^* equals a fixed number of sellers for a set of relatively large intervals. Lemma 14 and 15 characterize intervals where consumers continue sampling seller k only with some probability less than one. That is, only a fraction of consumers without a match samples seller k while all remaining consumers stop search.[11] Between k and K^*, no further radomized stopping occurs since all continuation values are strictly positive.

To see intuitively why random stopping occurs, consider the following argument. By lemma 4, the optimal price without tracking is a function of K^*, which, in turn, depends on the continuation value and thus on the level of search costs. As can be seen from equation (2.4) by substituting $p_j = p \; \forall \; j$, the continuation value V_{h+1} also depends on the no tracking price $p\big(K^*(s)\big)$, which in turn, depends on $K^*(s)$. That is, the continuation value under no tracking for any history h is given by $V_{h+1}\big(K^*, p\big(K^*(s)\big), s\big)$.

Ceteris paribus, a rise in search costs thus reduces the continuation value from search and leads to a lower search persistence K^*. However, if consumers sample fewer sellers in total, lemma 4 implies that the profit-maximizing price decreases, thus increasing consumer surplus. When search costs are such that a consumer is just indifferent between sampling seller k and stopping search, a marginal increase in s has only a marginal direct effect on the continuation value from search while the indirect effect through a lower k is large. Without random stopping, the discontinuity in K^* results in an inconsistency that rules out an equilibrium in pure strategies. Instead, consumers shift sellers' beliefs towards expecting more consumers with shorter search histories by sampling later sellers only with some probability less than one. This prevents price jumps and restores the equilibrium.

It remains to analyze the effect of search costs on prices in the absence of tracking. Considering only the intervals where consumers do not use mixed strategies, higher

[11] Recall that k identifies a seller's position in any consumer's search process and not a unique seller.

search costs imply a lower K^* by lemma 13 and thus a decrease in prices by lemma 4. When search costs are in a region where consumers follow a mixed stopping rule, prices must decrease as well. Consider an increase in search cost from $s = \hat{s}^*_K(K^*)$ to some $s > \hat{s}^*_K(K^*)$ requiring consumers to sample seller K^* only with some probability.[12] If search costs increase, the equilibrium probability of continuing to search decreases. This is because it reduces the mass of actively searching consumers with any history $h = K^* - 1$ such that sellers expect fewer consumers with long search histories and set lower prices. Hence, search persistence decreases smoothly in search costs.

Proposition 2.3 *Without tracking, the uniquely defined uniform price is weakly decreasing in search costs.*

The intuition behind this result follows immediately from lemma 4. Lower search costs increase consumers' search persistence, which reduces every seller's share of elastic demand and, thus, leads to higher prices. Proposition 2.3 provides a micro-founded theoretical explanation for some empirical papers suggesting that the internet does not always lead to lower prices despite reducing search costs. For example, Ellison and Ellison (2014) find that prices for used books are higher online than offline. In line with my model's predictions, the authors argue that higher prices obtain because sellers expect to sell mostly to consumers with high match values when consumers are willing to search longer due to lower search costs.

It is important to note that a mixed stopping rule that can continuously decrease a seller's belief about the average search history need not exist always. That is, consumers' search persistence may not decrease gradually but may immediately fall from $K^* > 1$ to zero. Then, the *no search* equilibrium, which always exists, is the only equilibrium. Formally, define by $\hat{s}_k(K^*) \in \{s : V_k\big(K^*, p\big(K^*(s)\big), s\big) = 0\}$ the threshold level of search costs such that sampling seller k conditional on $p(K^*)$ and K^* is worthwhile if and only if $s \leq \hat{s}_k$. Then,

Proposition 2.4 *If there exists an equilibrium with $K^* = \underline{K}^* \leq N$ for some level of search costs such that $\hat{s}_1(\underline{K}^*) < \hat{s}_k(\underline{K}^*)$ $\forall\ 1 < k \leq \underline{K}^*$, no consumer searches and the market breaks down if $s > \hat{s}_1(\underline{K}^*)$.*

Notably, the no search equilibrium prevails even though the surplus from search would be positive for consumers if they were able to commit to sampling less than a certain number of sellers. However, for any price making the first search worthwhile, the search persistence, resulting from consumer's sequentially optimal stopping decision, exceeds the search persistence for which the assumed price is maximizing profit. Instead, sellers anticipate consumers' search persistence conditional on initiating search

[12]The notation is explained in the appendix. For the argument however, it is sufficient to treat $\hat{s}^*_K(K^*)$ as some fixed threshold.

and want to set a higher price, making sampling even the first seller not worthwhile for consumers. Further, the conditions from proposition 2.4 rule out any mixed stopping rule as shifting sellers beliefs to a lower average search persistence is not feasible when the problem occurs at the first seller.

Proposition 2.4 applies if threshold search cost levels are not decreasing in the order of search. Importantly, this always happens for some level of search costs if the continuation value is not only decreasing with longer search histories. Technically, this depends on $g'(x)$ as well as $x \cdot g(x)$. In reality, a situation where the continuation value increases may in fact be quite common. Without knowing well what he is looking for, Bob might not be too enthusiastic about getting a new suit prior to searching. While sampling the first sellers however, Bob might learn that he likes particular kind of buttons and becomes excited about finding a suitable shirt. As a consequence, his interim continuation value from search might well exceed his expected surplus prior to search.

2.4 Comparative Statics of Tracking

As the subsequent analysis will show, the implications of tracking for overall welfare depend on the level of search costs and thus on consumers' search persistence as well. Because obtaining predictions that depend on the model's fundamentals requires additional structure, I first discuss how the effects of tracking vary with the search persistence parameter K^*.

2.4.1 General Analysis

The first result concerns equilibrium prices and immediately follows from the sellers' first order conditions.

Proposition 2.5 *For any $K^* > 1$, prices with and without tracking satisfy:*

$$p_1 < p(K^*) < p_{K^*}.$$

By proposition 2.5, the uniform no-tracking price exceeds the price consumers face at their first seller when searching with tracking but is strictly below the last seller's price. Moreover, a general result regarding profits is available when K^* under search with tracking is at least as high as under search without tracking:

Lemma 5 *If K^* is weakly larger under search with tracking, sellers' profits are strictly larger under search with than under search without tracking.*

The reason is fairly intuitive. If consumers sample the same number of sellers under both regimes, the aggregated distribution of types from all consumers is identical. Then, the only difference between tracking and no tracking is that in the former case, sellers can condition the optimal price on private information about consumers. By proposition 2.1, prices are increasing in the order of search under tracking and thus different from a uniform price. Hence, the uniform price is not profit-maximizing if better information is available, implying that tracking yields higher profits. Importantly, K^* is indeed weakly larger under search with tracking in many cases. If search costs are sufficiently low and the number of sellers not too large, $K^* = N$ irrespective of tracking. Consequently, it follows generally that sellers always benefit from tracking if there are not too many of them or search costs are sufficiently low.

Moreover, the market breakdown result stated in proposition 2.4 implies that if the continuation value from search is not decreasing in the consumer's search history, there exists a threshold search cost level $\hat{s}_1(\underline{K}^*)$ such that $K^* = 0$ for any $s > \hat{s}_1(\underline{K}^*)$ under search without tracking. Thus,

Proposition 2.6 *If there exists a $\underline{K}^* \leq N$ such that $\hat{s}_1(\underline{K}^*) < \hat{s}_k(\underline{K}^*)$ $\forall$ $1 < k \leq \underline{K}^*$, tracking leads to strictly higher consumer surplus and profits for $\hat{s}_1 > s > \hat{s}_1(\underline{K}^*)$.*

Proposition 2.6 holds irrespective of how much surplus sellers would extract from consumers via search history-based price discrimination. For $s \in [\hat{s}_1(\underline{K}^*), \hat{s}_1)$, it holds that $V_1(p_1^+, s) \geq 0$ under search with tracking. That is, the market under search with tracking is active at a level of search costs where the market without tracking is not. Consequently, tracking raises everyone's surplus. The reason is that tracking reduces the information asymmetry between consumers and sellers and thereby leads to sufficiently low prices, making initiating search worthwhile for consumers. In contrast, an equilibrium with low search persistence and low prices does not exist without tracking under the conditions of proposition 2.6. Proposition 2.6 also stands in contrast to the results derived by Zhou (2011). Without ex ante consumer heterogeneity, he finds that consumer search is inefficiently low when search is ordered and leads to a lower overall surplus.

As opposed to standard monopolistic group pricing, the following analysis suggests that tracking may raise overall consumer surplus even when there is no market breakdown in the absence of tracking. Similarly to group price discrimination, it has a market expansion effect since prices are lower for consumers with a lower expected willingness to pay.[13] However, discrimination based on search histories is likely to benefit consumers even more than standard group price discrimination. Rather than being part of mutually exclusive groups, consumers arriving at some seller k have already had the opportunity to buy at any seller $j < k$, implying that the different "groups" of consumers facing different prices are in fact subsets of one another. Moreover, the market price without tracking is above the first and below the last seller's price with tracking by proposition 2.5. Hence, there must exist a threshold history $\bar{h}$ such that only consumers with a history $h > \bar{h}$ pay discriminatory prices exceeding $p(K^*)$. Note that consumers with a niche taste search longer on average than consumers with a mass taste do. Hence, in expectation, some types are made better off from tracking while others are made worth off.

Corollary 2.1 *There exists a cut-off type $\tilde{x} \in X$ such that a consumer's expected surplus is reduced due to tracking if $x > \tilde{x}$.*

2.4.2 Linear Matching Probability

In this section, I impose further structure on the model to derive additional analytic results. In particular, I am interested in observing when the reduction of asymmetric information via tracking can lead to welfare and consumer surplus improvements, whether tracking always raises industry profit and, how these effects depend on search costs. Consider a linear matching probability function $g(x) = 1 - x$ and let the type x be uniformly distributed on $[0, 1]$. The total number of sellers is held constant at $N = 10$ but consumers may sample only $K^* \leq N$. Proposition 2.7 summarizes the findings from this section. Importantly, additional computations show that qualitatively identical results obtain from any linear matching probability function.

Proposition 2.7 *There exist two thresholds $s_1^p < s_2^p$ such that profits are strictly larger under search without tracking if $s \in (s_1^p, s_2^p)$. Besides, there exist two thresholds $s_1^w < s_2^w$ such that welfare is strictly larger without tracking if $s \in (s_1^w, s_2^w)$. Consumers surplus is always higher under tracking. The market breaks down without tracking for $s > s_2^w$.*

The prices under search with tracking can be obtained from solving the maximization problem as specified in (2.3). Under search without tracking, I observe that prices increase in K^*.

[13] In fact, Belleflamme and Peitz (2015) demonstrate how partitioning the demand into smaller intervals ("groups") has a non monotonic effect on consumer surplus due to two opposing factors. Due better information about the willingness to pay sellers charge lower prices from groups with a lower willingness to pay - eventually leading to an expansion of the market as consumers are served that would not have bought under the uniform price. However, as this information becomes more precise, the surplus left to each group is decreasing and approaches zero for infinitesimally small intervals.

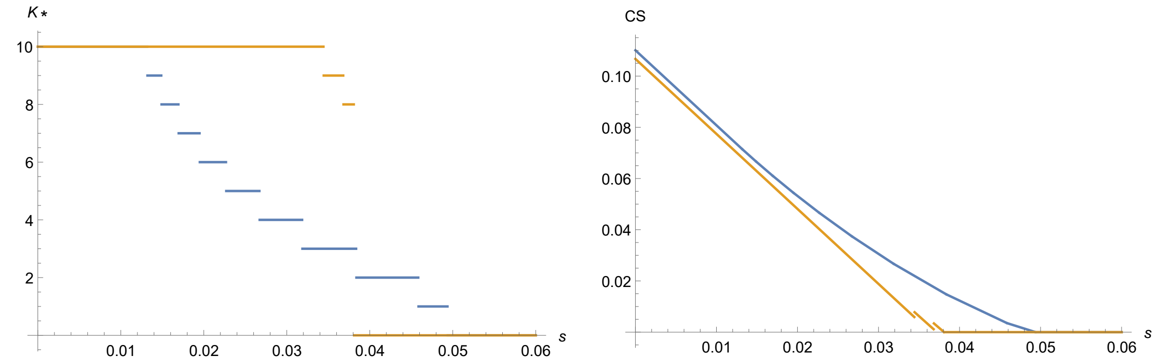

(a) Search persistence

(b) Consumer surplus

Figure a) shows the search persistence K^* as a function of search costs. Figure b) shows consumer surplus as a function of search costs. Orange represents search without tracking, blue search with tracking.

Figure 2.1: Search costs, search persistence and consumer surplus

The continuation value for every history and feasible K^* is given by equation (2.4). Under search with tracking where continuation values and threshold search cost levels are decreasing due to increasing prices, K^* can be derived from the set $\mathbb{K}$ as defined in equation (2.5). For search without tracking, the dependence of the continuation value on K^* via its effect on prices imposes an additional constraint on the optimal K^* as explained in section 3.3. Beyond $V_k(K^*, p(K^*)) \geq 0 \; \forall \; k < K^*$, it requires that $V_{K^*}(K^*, p(K^*)) \geq 0$ for the price that is optimal conditional on Kj. Lemma 13 specifies the search cost intervals for every $K^* \leq N$. Pure stopping strategies sometimes lead to dynamic inconsistencies due to the adverse effect of K^* on the price p. By lemma 14, consumers with the maximum possible search history in the market randomize over sampling the "last" seller for some levels of search costs. However, the range of the intervals where consumers would choose mixed strategies is relatively small compared to those where they choose pure strategies under the model's specifications. Therefore, I omit the calculation of the mixed strategies.[14]

The sharp drop of K^* under search without tracking at around $s = 0.038$ as displayed in figure (2.1a) illustrates the market breakdown result from proposition 2.4. If $s = 0.038$, the continuation value from sampling the first seller (expected surplus from search) falls below zero while to any other interim continuation value would still be positive. Thus, the threshold search cost is lowest prior to sampling search and the dynamic search inconsistency problem cannot be prevented by a randomized strategy. Hence, the market shuts down entirely and K^* equals zero. As a consequence, consumer

[14]The length of the intervals where consumers would mix can be inferred from figure (2.1b). The discontinuities between $s = 0.03$ and $s = 0.04$ would disappear if mixed strategies were accounted for. Importantly, neglecting the mixed stopping behavior does not affect the computation of consumer surplus, as $V_j = 0$ if consumers randomize over sampling j.

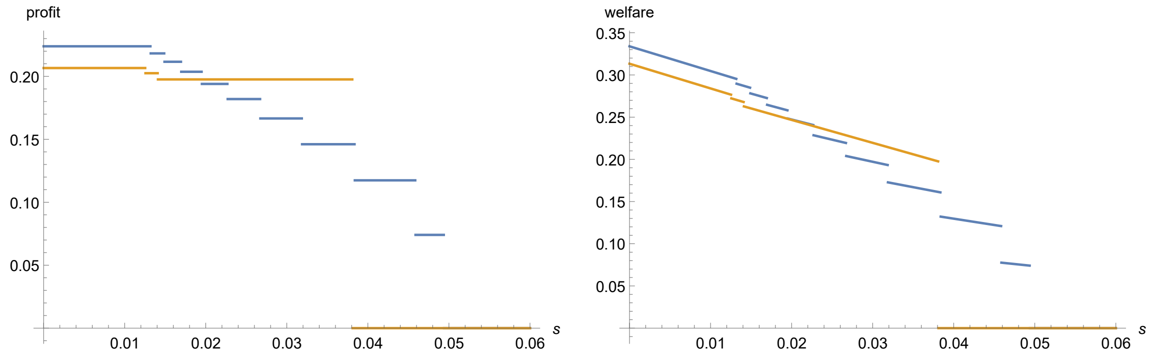

(a) Profits (b) Welfare

Figure a) shows total profits as a function of search costs. Figure b) shows welfare as a function of search costs. Orange represents search without tracking, blue search with tracking.

Figure 2.2: Search costs, profits and welfare

surplus is significantly lower under search without tracking when search costs hit the market-breakdown threshold as shown in figure (2.1b).

Moreover, it can be seen that tracking leads to higher consumer surplus even in the absence of the market breakdown. For search cost in the neighborhood of the cut-off level $s = 0.038$, this is not too surprising. Both consumer surplus with and without tracking are continuous functions of search costs. Since it is strictly higher with tracking at $s = 0.038$, it has to be higher for lower search costs as well. For the given linear specification of the model, tracking always raises consumer surplus. This holds irrespective of the fact that for some levels of search costs, consumers sample fewer sellers with tracking. The foregone consumer surplus from sampling the last sellers is, however, relatively small due to higher prices conditional on longer search histories. Thus, consumers always benefit on average.

Note that, surprisingly, tracking is not always maximizing industry profit as can be seen in figure (3.3a) for intermediate search costs. In particular, this happens if consumers' search persistence is significantly lower under search with tracking than without tracking. More precisely, the benefit of tracking for sellers comes from exploiting detailed information about consumers with long search histories. If search costs are too high however, consumers with long search histories (and high expected match values) anticipate that they will be left with almost no surplus and do not continue sampling additional sellers. While sellers would in general benefit from promising to leave a surplus to consumers compensating them for search costs, they cannot due to the hold-up problem.

Thus, the situation when facing a consumer with a long search history is comparable to Diamond (1971), where sellers have perfect information about a consumer's willingness to pay. In other words, the *Diamond Paradox* applies and consumers forego search, which mostly harms sellers who would have extracted most of the surplus. Figure (3.3a) shows that tracking raises sellers' profits otherwise and especially when the market would shut down without tracking for high search costs. Figure (3.3b) shows the effect of tracking on overall welfare, which depends on the level of search costs as well. There is a wide range of search costs for which foregone profits due to reduced search persistence cannot be offset by the increase in consumer surplus. For intermediate search costs, tracking thus reduces welfare. In fact, this is intuitive as the hold-up problem prevents the realization of matches especially from high value consumers.

2.5 Extension: Increasing Matching Probability

In markets where products are sufficiently complex and have a variety if differentiated features, the distinction between mass and niche consumers implying $g'(x) < 0$ seems a reasonable assumption. However, if products are more standardized, the major source of different match values may not lie in the nicheness of a consumer's taste. Rather, a consumer's match value will depend on her budget set. For instance, think of consumer electronics like flat screens. While not every brand's flat screen constitutes a match due to different preferences for diameter or energy consumption, the majority of consumers derive utility from all of its technical features. Consequently, their willingness to pay depends on mostly on the available income although they (still) buy the product only if the investigated features meet their individual preferences. Aguiar and Hurst (2007) find that low income consumers do less comparison shopping but still spend more time on shopping in total, suggesting that they spend more time than high income consumers on studying the products they buy. Given high income consumers inspect a product's features less carefully than low income consumers, they might be more likely to find a product suitable. In this section, I therefore analyze the case of a weakly increasing matching probability function $g(x)$ with $g'(x) \geq 0$.

Many of the arguments and results are either identical or simply a reversed version of those made in the previous section. In those cases, explanations and proofs are presented in an abbreviated form with complete references to the previous section. I first show under which conditions there exists an equilibrium that exhibits a decreasing price path. Second, I show that under those conditions, the equilibrium is unique. Moreover, I derive the equilibrium without tracking and compare market outcomes.

2.5.1 Search With Tracking

For $g'(x) \geq 0$, consumers with a high conditional match value are more likely to encounter a match early. Hence, intuition suggests that lower prices are charged from consumers with longer search histories in equilibrium. To construct such an equilibrium, suppose that consumers' expectations satisfy $p_1^e \geq p_2^e \geq ... \geq p_N^e$.

If consumers expect a decreasing price path, they might prefer to continue searching despite available matches at previous sellers. As the analysis of alternative expectations in the case of $g'(x) < 0$ has shown, this potentially leads to a new category of demand from consumers who might return to a match after sampling additional sellers. However, if potential gains from lower prices at forthcoming sellers are sufficiently smaller than search costs, continuing to search despite a match is not worthwhile. Then, consumers follow $\mathcal{R}^*$ and expected history-dependent demand can be characterized by equation (2.2) derived in section 3. In contrast to $g'(x) < 0$, demand from consumers with long search histories is more elastic if consumers follow $\mathcal{R}^*$ and types with a low conditional match value search longer on average. I thus obtain the opposite result of lemma 2.

Lemma 6 *Suppose that consumers always use the stopping rule $\mathcal{R}^*$. Then, the sequence of profit-maximizing prices is weakly decreasing and unique.*

As in the proof of lemma 2, uniqueness is due to log-concavity of the demand function, which is preserved under $\mathcal{R}^*$ for any seller along a consumer's search path. Lemma 6 does not yet complement proposition 2.1 for $g'(x) \geq 0$. The question remains under what conditions the continuation value from search after a successful match is always negative, thus rendering $\mathcal{R}^*$ indeed the optimal stopping rule?

Potential gains from lower prices depend on the changes in the elasticity of expected demand. Denote the decreasing sequence of optimal prices if consumers follow $\mathcal{R}^*$ by

$$\{p_k^*\}_{k=1,..,N},\ p_k^* \geq p_{k+1}^*\ \forall\ k < N.$$

Note that in any PBE where $\mathcal{R}^*$ is the optimal stopping rule, consumers' expectations must be correct and thus satisfy $p_k^e = p_k^*$. The stopping rule $\mathcal{R}^*$ is optimal always only if the continuation value from search conditional on an available match is weakly negative for all types and for all possible search histories. Extending previous notation, denote the continuation value conditional from sampling seller k conditional on an existing match and known type x_i by $V_k(x_i)$. Formally, $V_k(x_i) \leq 0\ \forall x_i \in X, k \leq N$ *iff*

$$g(x)\big(p_k^* - p_{k+1}^e\big) < s\ \forall x \in X,\ k \leq N. \tag{2.8}$$

Given any level of search costs $s > 0$, an upper bound $\hat{\Delta} > 0$ for the slope of $g(x)$ exists such that condition (2.8) holds for all matching probability functions with $g'(x) < \hat{\Delta}$. Consequently, $p_k^e = p_k^+ \; \forall k$ if $g'(x) < \hat{\Delta}$. Lemma 7 summarizes these findings.

Lemma 7 *There always exists a $\hat{\Delta} > 0$ such that under expectations satisfying $p_k^e = p_k^* \; \forall \; k \leq N$, the stopping rule $\mathcal{R}^*$ is optimal and $\{p_k^*\}_{k=1,..,N}$ constitute unique equilibrium prices for any matching probability function with $0 \leq g'(x) < \hat{\Delta} \; \forall \; x \in X$.*

The stopping rule $\mathcal{R}^*$ thus leads to the price sequence $\{p_k^*\}_{k=1,..,N}$, which implies that $p_k^e = p_k^*$, rendering $\mathcal{R}^*$ indeed optimal for $g'(x) < \hat{\Delta}$. However, other equilibria could exist for alternative consumer expectations, making some types adopt a stopping rule different from $\mathcal{R}^*$. However, it is possible to show that any expectations which are different from $\{p_k^*\}_{k=1,..,N}$ cannot constitute a PBE.[15] The steps towards this result are similar to those made in the previous section. First, observe that

Lemma 8 *In any PBE, expectations satisfy $p_k^e \geq p_k^* \; \forall \; k \leq N$ if $\hat{\Delta} > g'(x)$.*

Intuitively, the statement holds for the following reason. Begin with the last sellers whose price is expected to be below p_k^*, i.e. $p_k^e < p_k^*$. If despite alternative expectations, consumers' stopping behavior at previous sellers remains unchanged, seller k would maximize profits by deviating to $p_k = p_k^* > p_k^e$ as this price is profit-maximizing conditional on $\mathcal{R}^*$.

If instead, consumers' stopping behavior changes due to alternative expectations, it can affect only expected demand types with $x > \bar{x}$ for seller k, where $\bar{x} > p_k^e$. The threshold $\bar{x}$ characterizes the type whose matching probability $g(\bar{x})$ is too low to make sampling seller k despite available matches worthwhile. Consequently, expected demand from types $x_i > \bar{x}$ can at most in- but not decrease compared to demand that arises under $\mathcal{R}^*$. Since raising the price to $p_k^* > p_k^e$ leads to higher profits even in the absence of this extra demand by lemma 6, setting a price higher than p_k^e must constitute a profitable deviation when demand from types $x_i > \bar{x} > p_k^e$ as well. The same argument can be applied to any seller k' whose expected price is supposed to satisfy $p_{k'}^e < p_{k'}^*$, leading to the above lemma. As in the analysis in section 3, the lower bound on consumers' expectations leads to equilibrium uniqueness.

Proposition 2.8 *In any PBE, expectations must satisfy $p_k^e = p_k^* \; \forall \; k \leq K^*$. Hence, the equilibrium with the increasing price path characterized by $\{p_k^*\}_{k=1,..,N}$ is unique.*

The proof proceeds along the same lines used in proposition 2.1. As the first seller knows that prices from all forthcoming sellers are expected to be higher than his own

[15] As it is the case in most search models, there always exists an equilibrium in which consumers expect arbitrarily high prices and no consumer searches.

price if he charges a price in the neighborhood of p_1^*, he has local monopoly power over its demand because consumers would follow $\mathcal{R}^*$. Since p_1^* is the profit-maximizing price under $\mathcal{R}^*$, the first seller will always find it optimal to set a price equal to p_1^+. Given that the first seller sets $p_1 = p_1^*$, seller 2 can makes consumers follow $\mathcal{R}^*$ (and thus obtain local monopoly power) by setting $p_2 = p_2^*$ for the same reasoning. By repeatedly applying this argument for all forthcoming sellers, one obtains the above result.

Search persistence. As in the previous analysis, consumers sample up to K^* sellers if $V_{K^*}\big(s\big) \geq 0$ and $V_k\big(s\big) \geq 0 \; \forall \; k < K^*$ given equilibrium prices. In fact, the fact that prices are decreasing instead of increasing does not change the computation of K^*. That is, $K^* \in \mathbb{K}$, where $\mathbb{K}$ is given by (2.5). Since $\hat{K}(s') \subseteq \hat{K}(s)$ for $s' > s$ holds as well, K^* is decreasing in s, resulting in a set of disjoint search costs intervals for different K^*.

If continuation values are not always decreasing for higher search histories, there might be jumps in K^* such that search persistence decreases by more than one seller at some search cost threshold. Using previous notation, let $\hat{s}_k(K) \in \{s : V_k(K, s) = 0\}$ define the search cost threshold above which sampling k is not worthwhile for a given search persistence K. For non-decreasing continuation values, there exists a K such that $\hat{s}_k(K) \neq \underline{\hat{s}}(K) \in \min\{\hat{s}_k(K)\}_{k=1,..,K}$. Then, $K^* = K$ only if $s \leq \underline{\hat{s}}(K)$ and $K^* = j$ for $\hat{s}_{j'}(j) \geq s > \hat{s}_j(K)$ where seller j's threshold equals $\hat{s}_j(K) = \underline{\hat{s}}(K)$. The next threshold $\hat{s}_{j'}(j)$ is obtained from $\hat{s}_{j'}(j) \in \min\{\hat{s}_k(j)\}_{k=1,..,Kj}$ and specifies that $K^* = j'$ for $s > \hat{s}_{j'}(j)$ and so forth.

2.5.2 Search Without Tracking

Deriving uniform price equilibrium and its uniqueness under search without tracking is identical to the case of $g'(x) < 0$. The reason is that due to the uniform price, consumers follow the same stopping rule $\mathcal{R}^*$. Hence, the seller's problem is completely characterized by (2.7) with expected demand given by equation (2.2). Thus:

Lemma 9 *A unique uniform price equilibrium exists under random search.*

The only difference compared $g'(x) < 0$ is that demand is now more elastic for higher degrees of search persistence. That is:

Proposition 2.9 *Let $p(K^*)$ be the unit random search price if consumers maximum willingness to search is K^*. Then,*

$$p(K_2^*) < p(K_1^*) \text{ if } g'(x) > 0 \text{ and } K_2^* > K_1^*$$

Notice that the intuition for the result is a simple reversion of the statement before. A higher search persistence by consumers increases the probability that a consumer has a long history. Since probabilities for all feasible search histories must add up to one, sellers put less weight consumers with short histories and more weight on consumers with long search histories if K^* increases. Since demand from consumers with longer search histories is more price elastic for $g'(x) \geq 0$, the profit maximizing price decreases.

Search persistence. For a constant price, the continuation value from search decreases for longer search histories. This is because a consumer's expected type $\mathbb{E}_h[x]$ decreases in h and so does the instantaneous expected surplus from the next seller, $g(x)(x-p)$, if $g'(x) \geq 0$. Since consumers become increasingly pessimistic for higher search histories, continuation values are decreasing in h, implying that K^* decreases gradually. That is, consumers sample $K^* = K$ sellers if $s > \hat{s}(K)_K$ where $s_K(K) \in \{s : V_K(p(K), s) = 0\}$. Since $p(K)$ is decreasing in K, dynamic search inconsistencies as in the case of $g'(x) < 0$ cannot emerge. Hence, the market remains active for all $s < s_1(1)$. Moreover, consumers never have to choose mixed stopping rules.

2.5.3 Comparative Statics of Tracking

Since search persistence decreases gradually with search costs and $V_1(p(1), s) = V_1(p_1, s)$, it follows that consumers search both under tracking and no tracking *iff* $s < \hat{s}_1(1)$. Contrary to the case of $g'(x) < 0$, there thus exists no level of search costs for which tracking must lead to higher profits and consumer surplus due to a market breakdown without tracking. The comparison of prices is immediate after switching the sign of $g'(x)$ in the proof of proposition 2.5:

Corollary 2.2 *For any search persistence $K^* > 1$,*

$$p_1 \geq p(K^*) \geq p_{K^*}.$$

with strict inequality if $g'(x) > 0$.

By corollary 2.2, sellers maximize profits by charging search history dependent prices, which differ from the uniform price without tracking. That is, reduced asymmetric information due to tracking enables sellers to extract more expected surplus from a consumer with a particular history h. Unless $g'(x) = 0$, I obtain:

Lemma 10 *If K^* is weakly higher under tracking than under no tracking, sellers' profits are strictly larger under search with than under search without tracking.*

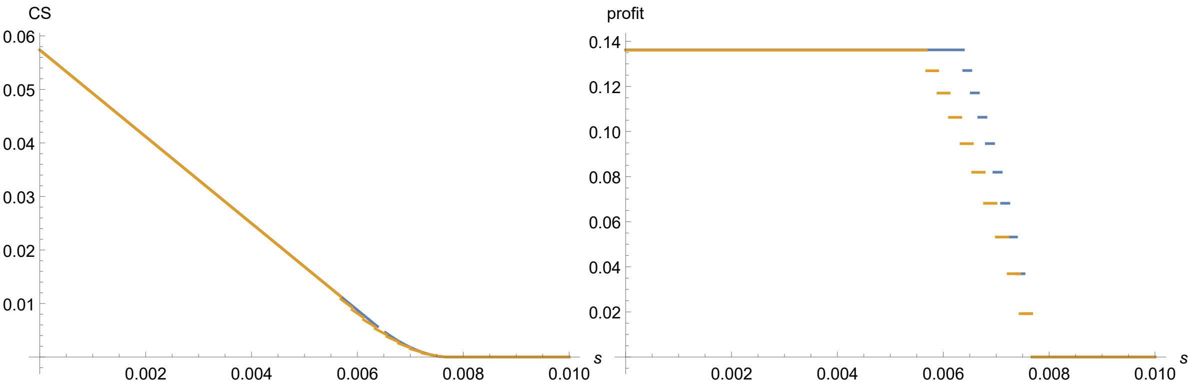

(a) Consumer surplus (b) Profits

Figure a) shows consumer surplus as a function of search costs. Figure b) shows total profits as a function of search costs. Orange represents search without tracking, blue search with tracking.

Figure 2.3: Search costs, consumer surplus and profits

To obtain specific results on the effect of search costs, I again impose additional structure. Consider a linear matching probability function$g(x) = 0.1x$, which is increasing in x. As before, let the type x be uniformly distributed on $[0, 1]$.[16] The total number of sellers is held constant at $N = 10$ but consumers may sample only $K^* \leq N$.

The computation proceeds as follows. First, I compute prices under search with tracking for any history h and under search without tracking for any search persistence K^*. Second, by calculating continuation values for every possible history, I obtain the optimal search persistence K^* as described in section 5.1 and 5.2 for every level of search costs. Recall that the decreasing price path under search with tracking may not constitute an equilibrium if search costs are too low or potential gains from price savings too high. While not displayed here, the maximum price difference a consumer can expect amounts to roughly $\Delta_p = 0.005$. Since the highest matching probability a consumer might have equals 0.1, condition 2.8, which ensures that a unique equilibrium exists, is satisfied for $s > 0.0005$. Figure (3.4a) shows that tracking hardly affects consumer surplus. That is, gains for longer searching consumers from lower prices are offset by higher starting prices for all consumers. In contrast, profits can be much larger due to tracking as shown in figure (3.4b). In fact, this is because K^* decreases faster if search is without tracking. While under tracking, consumers with long search histories may find continuing to search worthwhile such that match values are realized, the price without tracking may prevent them from search. As the surplus left to consumers even if they continue searching is small, there is only a marginal effect on consumer surplus.

[16]The slope is chosen to be small to ensure equilibrium existence.

2.6 Application: Endogenous Tracking

In this section, I apply the consumer search framework to study whether tracking arises endogenously. For this purpose, I consider a consumer's dynamic choice about preventing tracking. Since the processing of personal data seems to be the default on the internet for its use goes far beyond price discrimination only, a seller's choice about tracking is only about deciding on whether to use the available data for price discrimination. Note that this implies a seller decides about using tracking at the price setting stage. Hence, the problem is simply part of his profit maximization. More precisely, if anything but a search history-independent price is optimal, it must hold that using tracking is the dominant strategy.

To model a consumer's tracking choice, I assume that every time before she samples a seller, she can either disclose her entire search history, i.e. allow tracking (T) or not disclose her search history and thus not allow tracking (NT). Search histories contain the number of all sellers previously visited, independently of whether the consumer had chosen NT or T when sampling previous sellers.[17] In reality, consumers usually also have the option to erase their histories, for example by deleting cookies. In the discussion section, I show that the predictions I obtain are robust to this extension if sellers can distinguish between a consumer who deletes her cookies and a consumer who just started searching for a particular product. Briefly, this assumption is motivated by the observation that by deleting cookies, consumers erase their entire search profile. However, a consumer who has not deleted cookies but just started searching for a product should still have an "unrelated" search history. The signal sent to sellers when deleting cookies is then identical to choosing NT and, thus, is redundant. Denote a consumer's tracking decision by $d \in \{T, NT\}$.

Timing. At first, sellers set prices for every possible search history and choice of d and nature draws each consumer's type. Consumers search by sampling sellers sequentially at a cost s per seller. Prior to each search attempt, consumers choose d. Sellers observes a consumer's search history h if $d = T$ and nothing but $d = NT$ otherwise. The consumer observe the price conditional on her disclosure strategy and search history as well as her match utility. Lastly, she decides whether to buy, to return to a previous seller, to continue or to stop search.

[17] Web-browsing with the "do not track" request option, where cookies are still stored on the consumer's device but simply not processed, fits this assumption fairly well. However, results are robust to the modified assumption that search histories contain only the number of sellers for which tracking had been enabled. This is because the search history a consumer can disclose would still be weakly increasing in the number of sellers and thus not affect the price path with tracking.

Equilibrium concept. As before, the equilibrium concept is Perfect Bayesian Nash equilibrium. Extending the strategy space by the choice of d implies additional PBE conditions which are not present in the baseline mode. First, consumers choose d in order to maximize their expected surplus. Second, sellers' beliefs about a consumer's search history must be consistent with the consumers' disclosure strategy. Third, consumer's beliefs about prices must be consistent with sellers' equilibrium pricing strategies, and, thus, with their own disclosure strategy.

I analyze the cases of $g'(x) < 0$ and $g'(x) \geq 0$ separately and begin with $g'(x) < 0$. Denote by $p(NT)$ the uniform price set by sellers upon observing the choice NT. For prices under tracking, use the previous notation p_k with the index indicating the seller's position in the consumer's search process. Note that the choice about $d \in \{T, NT\}$ is without commitment and only affects the information revealed to the next seller while the prices from additional sellers remain unaffected. Hence, the consumer's decision to disclose her history is purely myopic as it only depends on the difference between the next price she can expect under tracking p^e_{h+1} and no tracking $p^e(NT)$. If $p^e_k = p^e(NT)$ for some consumer with history $h = k - 1$, I impose that a consumer stays with the default option, which is search with tracking. This tie-breaking rule is without loss of generality as the alternative rule would imply the same equilibrium outcome. Since consumers choose $d = T$ prior to sampling seller k if only if $p^e(NT) \geq p_k$, it is sufficient to restrict attention to single cut-off strategies, where such a strategy is defined as follows:

$$\textit{Choose } NT \;\; \forall \; h \geq \hat{h} \in \mathbb{N} \tag{2.9}$$

For brevity, denote the above defined single cut-off strategy by $\hat{h} \leq K^*$. To see why this restriction does not constrain equilibrium strategies, suppose that $p^e(NT) < p^e_k$ such that a consumer chooses $d = NT$ at seller k. Since her search history at any forthcoming seller will be $h > k$, forthcoming sellers' prices always satisfy $p^e(NT) < p^e_j \; \forall \; j > k$. Consequently, $d = NT$ must be optimal for all $h \geq k$ if it is optimal at k.

In any PBE, sellers anticipate the equilibrium strategy $\hat{h}$ and can thus condition $p(NT)$ on $h \geq \hat{h}$ when observing $d = NT$. Denote the profit-maximizing price conditional on the cut-off strategy $\hat{h}$ by $p(\hat{h})$. Then, $p(NT) = p(\hat{h})$ in equilibrium. Recall that by previous notation, $p_{\hat{h}+1}$ denotes the price set by a seller who observes a browsing history $\hat{h}$ and thus knows that his position in the consumer's search process is $\hat{h}+1$. Without imposing equilibrium strategies yet, the following lemma compares prices with tracking and without tracking for arbitrary single cut-off strategies $\hat{h}$.

Lemma 11 *Given $\hat{h}$, there always exists a unique optimal price $p(\hat{h})$ with*

$$p(\hat{h}) > p_{\hat{h}+1} \ \forall \ \hat{h} < K^* - 1,$$

and $p(\hat{h}) = p_{\hat{h}+1}$ for $\hat{h} = K^*$.

The intuition behind lemma 11 is the following. When hiding the search history for all histories $h \geq \hat{h}$, sellers observing $d = NT$ attach positive probabilities on all histories $h \geq \hat{h}$ and zero probability on $h < \hat{h}$. Most importantly, $\mathbb{P}\big(h > \hat{h}\big) > 0 \ \forall \ \hat{h} < K^* - 1$, and hence the optimal price conditional on observing $d = NT$ is chosen with respect to a weighted demand function that is always less elastic than expected demand from a consumer disclosing $h = \hat{h}$ (implying $\mathbb{P}(h = \hat{h}) = 1$).

By lemma 11, a consumer whose search history equals the cut-off history is charged a lower price if she allows tracking ($d = T$) than if she does not ($d = NT$). However, by construction, a consumer with a history of $h = \hat{h}$ chooses $d = NT$, implying that profitable deviations exist at least for some consumers. The uniqueness result in the following proposition is an immediate consequence of this contradiction.

Proposition 2.10 *There always exists a unique PBE with the disclosure strategy $\hat{h} = K^*$ and a conditional no tracking-price $p(NT) = p_{K^*}$.*

Proposition 2.10 states that the search history is always disclosed in the unique equilibrium, leading to unrestricted tracking and price discrimination. Existence can be shown by means of an example. Simply consider an equilibrium with search history-based prices $p_1 < p_2 < ... < p_{K^*}$, a disclosure strategy $\hat{h} = K^*$ and a no tracking-price $p(NT) = p_{K^*}$. Since K^* is the maximum number of sellers a consumer is willing to sample, $\hat{h} = K^*$ means that no actively searching consumer chooses $d = NT$ and sellers should never observe NT. Denote by $\mu(h) \in [0, 1]$ a seller's out-of-equilibrium belief that the search history of a consumer having chosen NT equals h. Suppose it satisfies $\mu(h) = 0 \ \forall h < \hat{h}$ such that $\mu(K^*) = 1$. Since p_{K^*} maximizes profits conditional on $\mu(K^*) = 1$, sellers have no incentive to deviate from $p(NT) = p_{K^*}$. Besides, consumers have no incentive to deviate to another disclosure strategy $\hat{h}' < \hat{h}$ since $p(NT) \geq p_k \ \forall \ k \leq K^*$.

The uniqueness result is based on an unravveling mechanism similar to Milgrom and Roberts (1986). For any alternative cut-off strategy $\hat{h} < K^*$, sellers' beliefs must satisfy $\mu(h) = 0 \ \forall \ h < \hat{h}$. Hence, the optimal price conditional on observing NT satisfies $p(NT) = p(\hat{h}) > p_{\hat{h}+1}$ by lemma 11. Since consumers with a search history $h = \hat{h}$ can obtain the price $p_{\hat{h}+1}$ by allowing tracking prior to sampling seller $k = \hat{h}+1$, they always have an incentive to deviate from any cut-off strategy $\hat{h} < K^*$.

2.6.1 Increasing Matching Probability

If $g'(x) \geq 0$ but not too large, prices are monotone decreasing in search histories as shown in section 2.5. Thus, it follows that the optimal disclosure strategy belongs to the set of single cut-off strategies as well. However, the reverse pricing pattern requires to slightly adjust the notion of single-cut-off strategies, abbreviated by $\check{h}$:

$$\textit{Choose } NT \ \ \forall\, h \leq \check{h} \tag{2.10}$$

Denote by $p(\check{h})$ the seller's optimal price conditional NT and the consumer cut-off strategy $\check{h}$. Analogously to lemma 11, one can show that:

Lemma 12 *Given $\check{h}$, there always exists a unique optimal price $p(\check{h})$ which satisfies*

$$p(\check{h}) > p_{\check{h}+1} \ \forall\, \check{h} > 0,$$

and $p(\check{h}) = p_{\check{h}+1}$ for $\check{h} = 0$.

The distinction between a a search history h and the corresponding position $h+1$ in the search process of a seller observing h is again crucial to understand the implications of lemma 12. For illustration, suppose that $\check{h} = 1$ implying that consumers choose $d = NT$ if $h \in \{0, 1\}$. In any PBE, sellers would know that $h \in \{0, 1\}$ if $d = NT$ and $h \ngeq 1$. By lemma 12, the resulting optimal price $p(NT)$ then exceeds p_1, the price sellers would set if they knew that $h = 1$ with certainty. However, any consumer with a history of $h = 1$ can choose $d = T$ and convey their history to sellers, thereby obtaining a lower price than $p(NT)$. The example shows that $\check{h} = 1$ cannot be an equilibrium strategy for consumers. In fact, the unraveling argument applies again for any cut-off strategy $\check{h} > 0$ such that complete tracking remains as the unique equilibrium:

Proposition 2.11 *There always exists a unique PBE with the disclosure strategy $\check{h} = 0$ and a conditional no tracking-price $p(NT) = p_1$.*

While the unraveling mechanism illustrated in the above example is reversed compared to proposition 2.10, the proof of proposition 2.11 is otherwise identical. For any cut-off strategy $\check{h} > 0$ where NT is chosen from consumers with strictly positive search histories, there always exists consumers who are better off from tracking even though their history belongs to $h \leq \check{h}$. Hence, only $\check{h} = 0$ constitutes an equilibrium strategy.

The major implication of propositions 2.10 and 2.11 is that there is no privacy in the market because consumers rationally approve tracking at all stages during the search process. Since the consumers' best-response to price discrimination is a myopic decision based merely on the very next seller's price, the equilibrium outcome need not be consumer surplus maximizing.

2.7 Discussion

In this section, I discuss the robustness of my findings via several extensions to the model. Some extensions constitute an entirely new model for future research which cannot be discussed in every detail here. In this case, I only provide a brief discussion of what changes to expect. Lastly, I explain how the model contributes to the discussion of whether random search can be stable (Armstrong, 2017).

No commitment

Since consumers search with free recall, sellers might not only be interested in discriminating between consumers, but also to discriminate a consumer based on whether she arrives for the first time or whether she is returning. Indeed, Armstrong and Zhou (2010) focus on this issue. In my baseline model, the timing does not allow for within consumer price discrimination. Instead, I implicitely assume that sellers can commit to a price they will charge from returning consumers under the constraint that this price is equal to the price charged at the first encounter. Indeed, this is not without loss of generality. Due to free recall, a multiplicity of equlibria might arise without commitment since consumers could form any beliefs about the return-price. However, the commitment assumption could be relaxed by introducing a small but positive cost ϵ for returning to a previous seller. This is because when a consumer with history h decides to return to some seller $k < h$, she does so only if she expects $v_{ik} > p_k^R + \epsilon$, where p_k^R is the expected return-price. Hence, sellers have an incentive to raise their price at least by ϵ, making returning not worthwhile for some consumers. This argument holds for any $p_k^R < \bar{v}$ such that the arising hold-up problem prevents any consumer from returning if sellers have no commitment power. The reason why the *Diamond Paradox* arises in the market for returning consumers but not in the market for "fresh" consumers is because only in the former, consumers already know their willingness to pay, rendering their return decision informative for sellers.[18]

[18] This point is also made in Armstrong and Zhou (2010)

Heterogeneous search costs

Another dimension in which consumers naturally differ from one another might be the individual search cost. Should we expect countervailing effects from introducing search cost heterogeneity into the current framework, mitigating the search dynamics and welfare implications derived? Heterogeneous search costs imply that consumers differ with respect to their search persistence K^*. Under search with tracking, prices do not depend on K^*. Thus, there are no new qualitative effects of search cost heterogeneity, despite leading to heterogeneous stopping by consumers.

Under search without tracking, heterogeneous search costs mitigate the dynamic search inconsistency problem. Since for any marginal increase in search costs, there is only a marginal consumer reducing her search persistence, there are no jumps in K^* and prices react smoothly to changes in search costs. That is, there is again no effect of search cost heterogeneity except for making mixed stopping rules disappear. The inefficiency problem due to market inactivity as stated in proposition 2.6 also persists under search cost heterogeneity. While indeed some consumer will always search for reasonable levels of search costs, those with low search costs search particularly long, thereby driving up the price without tracking. Hence, consumers with high search costs abstain from search entirely, driving up prices even more and exlcuding additional consumers from participating in search.

Deleting cookies

By deleting cookies. consumers might be able to reset their their seacrh history to $h = 0$ and thereby trick sellers. However, it is possible to show that deleting cookies and not allowing tracking are equivalent with respect to their signal under some mild assumptions. These are: (1) cookies saved on a device cannot be manipulated but only erased completely or not at all. While a minority of internet users might be capable of deleting only particular cookies from their computer, the majority of users is restricted to the standard software which typically enforces this all or nothing property. And (2), sellers can distinguish between a consumer who deleted all her cookies and a consumer who only began searching for a product.

Note also that this second assumption can be derived from more fundamental assumptions about online browsing. A consumer who only begins searching for a particular product still has a non-empty browsing history including search for other product categories and various online services. Denote this “extended” browsing history by $\mathcal{H}$.

The browsing history $\mathcal{H}$ is empty only if a consumer deletes her cookies and otherwise satisfies $\mathcal{H} \notin \emptyset$. In addition, note that $h = 0 \not\Rightarrow \mathcal{H} \in \emptyset$.

Denote the choice of deleting cookies by $t \in \{KC, DC\}$, where DC denotes the choice to delete cookies while K refers to keeping them. Focusing on sellers' belief about a consumer's (history-) type induced by the signal $t = DC$, it turns out there is no difference to the signal $d = NT$. Under NT, sellers cannot observe h. Under $t = DC$, sellers observe $h = 0$ but know that the consumer has deleted cookies since $\mathcal{H} \in \emptyset$. Hence, they know that they do not know the true h, which is equivalent to not knowing h at all.

Denote by $\{d, t\}$ the consumer's action tuple regarding the decision to allow tracking and keep her cookies. Due to identical signaling effects, it holds that sellers' beliefs satisfy:

$$\mu(h|T \wedge DH) = \mu\big(h|NT \wedge DH) = \mu(h|NT \wedge KH) \;\forall\; h.$$

Since the sellers' beliefs determine prices, these actions are thus payoff-equivalent from the perspective of consumers as well.

Now suppose that a fraction of consumers indeed searches with a new device and that her browsing history cannot be distinguished from a consumer who deletes her cookies. Equivalently, there might be a fraction of consumers who always disable tracking by default because they have a strong preference for privacy. If the fraction of those consumers is sufficiently large, the equilibrium from proposition 2.10 cannot be sustained since sellers must attach positive probability on shorter search histories conditional on observing NT (or equivalently, DC).

Notably, the first order condition of a seller observing NT shows that the optimal no tracking price $p(NT)$ is weakly decreasing in the cut-off strategy $\hat{h}$.[19] Besides, $p_{\hat{h}+1}$, the price a consumer with a history $h = \hat{h}$ would obtain if she chose T, is increasing as shown by lemma 2. Denote by α the fraction of consumers who disable tracking by default/ begin search with new devices. It follows that:

Corollary 2.3 *For $\alpha > 0$ but not too large, there exists a unique PBE with the disclosure strategy $0 < \hat{h} \leq K^*$ and the no tracking price $p(NT)$ solves*

$$p(NT) = \arg\max_{p} \big((1-\alpha) \sum_{k=\hat{h}+1}^{K^*} \phi_k D_k(p) + \alpha \sum_{k=1}^{K^*} \phi_k D_k(p)\big)p.$$

[19]See the proof of lemma 11

Stable random search

According to Armstrong (2017) random search in the WAR model is unstable. He argues that random search from the part of consumers depends crucially on consumers' expectations about other consumers' perfectly random search behavior. If instead, one seller S becomes more salient than its competitors such that both consumers and sellers should expect that an arbitrarily small but positive mass of consumers is more likely to sample S first, S will optimally set a lower price than its competitors, as shown by Armstrong and Zhou (2011) or Zhou (2011). If consumers are free to choose at what seller to begin searching while all sellers' products are ex ante identical, this creates an incentive for all consumers to begin searching at seller S. Consequently, the seller sampled first is not a random choice and search becomes partially ordered. Similarly, consumer beliefs about which seller to sample afterwards can tip easily and so can beliefs about all sellers' position in the search process. Then, random search becomes perfectly ordered as the same argument applies to the second seller etc.

The market tipping property of random search prevails because sellers who are visited first have an incentive to price lower than sellers who are sampled later in the search process. In the current framework instead, this is not true if $g(x)$ is weakly increasing, suggesting that random search would be stable. Suppose that a positive mass of consumers would not search randomly but begin with a particular seller M such that M has a larger share of consumers with shorter search histories. As $g(x)$ is increasing and high types are less likely to be among consumers with longer search histories, seller M has an incentive to raise its price above the uniform price charged by others.[20] Hence, the more salient seller M sets the highest price if $g'(x) > 0$. This prevents the search market from tipping as due to its higher price, consumers would avoid searching for seller M first, including those that were expected to sample it first on purpose. Consequently, consumers prefer to search randomly over coordinating on more salient sellers.

2.8 Conclusion

Most of users' activities across the internet are tracked by third parties. Accessing these browsing data is particularly attractive to sellers if the average search behavior varies across different consumer types. Then, tracking a consumer's search path enables sellers to learn about a consumer's willingness to pay and gives rise to search path-

[20] By not raising the price by too much, consumer's stopping rule would remain unaffected so that the stopping rule remains unaffected at least for small price deviations.

dependent price discrimination. This paper presents a rich and tractable framework integrating consumer heterogeneity and tracking into a model of sequential consumers search to address the major implications of tracking for market outcomes.

First, I show that, in the unique equilibrium, tracking implies search history-dependent pricing. Specifically, prices increase in the order of search if the difference between consumer types is the peculiarity and nicheness of their taste. Since niche consumers are more likely to search longer, demand from consumers with long search histories is less elastic. Consequently, sellers set higher prices conditional on observing longer search histories.

Second, I compare the market outcome under search with tracking with the unique equilibrium under search without tracking to evaluate its welfare consequences. Because initial prices are lower while later prices are higher than the price under search without tracking, the surplus of niche consumers decreases while the surplus of mass consumers increases due to tracking. Besides, overall welfare effects depend how tracking affects consumers' search persistence, which, in turn, depends on the level of search costs. For a wide range of intermediate search costs, consumers sample more sellers in the absence of tracking because the average no-tracking price makes continuing search with a long search history more attractive. This may cause welfare losses due to forgone matches, particularly at the cost of sellers who would have extracted most of the matching surplus. However, tracking increases welfare if search costs are very low such that search persistence remains unaffected and if search costs are very high. In the latter case, this is because low prices conditional on short search histories ensure that consumers have an incentive to begin searching, thus keeping the market active. In the absence of tracking instead, the no-tracking price is often too high, thus leading to a market breakdown for the same level of search costs. Perhaps surprisingly, consumers may always be better off from tracking whereas sellers make less profits at least for some search costs.

Third, I investigate whether tracking prevails endogenously when consumers can dynamically opt out from tracking. I find that, in the unique equilibrium, consumers always prefer to disclose their search history as the price conditional on hiding it always exceeds the price at least some consumers could obtain after disclosing it. Therefore, the entire search history is disclosed in equilibrium. Even though the equilibrium is unique, the full tracking prediction is not cast in stone. As discussed in the previous section, partial tracking, where consumers disclose their search history only up to some threshold, obtains if a positive mass of consumers always disables tracking by default. Besides, the endogenous tracking result is interesting not only because it explains why many internet users do not prevent tracking, but also because it has important implications for policy makers. That is, if tracking is surplus-increasing, no intervention is necessary because tracking prevails even though it entails increasing prices.

Often, sellers refrain from personalized pricing because of consumers' prejudices against price discrimination or legal uncertainty. Then, improved targeting constitutes an alternative practice to capitalize on traceable search histories. That is, a seller possessing multiple products of the same category might be able to use information conveyed through search histories to offer more suitable products to individual consumers. Following Johnson and Myatt (2006), a seller's product choice could be integrated in this paper's framework by allowing sellers to rotate the matching probability function. Again, it would be interesting to examine whether the profit-maximizing design induces inefficiently low search persistence or whether it can be welfare increasing as well.

Finally, Turow et al. (2009) find that the majority of consumers oppose personalized pricing, thus confirming the anecdotal evidence that significant prejudices against tracking still prevail. In the light of the fact that tracking often has desirable consequences for consumer surplus and welfare, the question of where this negative view comes from deserves more attention. If consumers obtain additional utility from anonymity, avoidance of tracking and personalized pricing may, of course, be welfare maximizing despite the foregone matching surplus. If, however, the preference for anonymity is based on false beliefs about the consequences of tracking, consumers may be harmed from being misinformed.

3 Heterogeneous Fairness Views[1]

3.1 Introduction

Distributional preferences are important inputs into social decision making, including taxation, redistribution and charitable giving. Different studies suggest that distributional preferences depend on people's fairness ideals. Many people are either egalitarians or libertarians, meaning that they either support or oppose redistribution always, irrespective of what they believe about the source of the inequality. The remaining fraction may be considered as liberal egalitarians, who are in favor of some redistribution unless they hold the poor fully responsible for their economic situation. Often however, only imperfect information is available about the source of economic success or failure. As a consequence, support for redistribution not only depends on one's fairness ideal, but also on what is believed about the factors causing economic inequality. In the light of this multiplicity of factors driving people's fairness views, it is not too surprising that many of the most frequently debated topics in everyday politics center around inequality and support for redistribution. Yet, despite the apparent heterogeneity in distributional preferences, our understanding of their premises, people's fairness ideals and beliefs about the accountability for economic success, is still incomplete. Because support for redistributive policies can neither be reasonably predicted nor fully understood without knowing why people have heterogeneous preferences, advancing our knowledge about what shapes these premises is a prerequisite for effective policy making. In this paper, I explore one possible explanation for the heterogeneity in the afore-mentioned dimensions, which is variation in the experience of economic success and failure.

Most importantly, I ask whether experience leads to systematically different beliefs about the influence of luck or effort when facing uncertainty about the source of inequality. To address this question, I use an experimental design that allows me to

[1]I am very grateful to Henrik Orzen, Dietmar Fehr, Hans-Peter Grüner, Justin Leduc, Nikos Nikiforakis, Jean-Robert Tyran and Stephan Meier for helpful comments and valuable suggestions.

observe distributive preferences under uncertainty and, more importantly, relate those preferences to participants' beliefs by comparing choices with those made under certainty about the earnings procedure. Additionally, I investigate whether experience affects the adoption of different fairness ideals by measuring distributive preferences while leaving no room for reasonably different beliefs about the role of luck.

In any natural environment, experience is endogenous and depends on a large number of unobservables. Moreover, any natural environment exhibits some degree of uncertainty about the process leading to inequality. However, relating people's distributive choices under uncertainty to their beliefs about the role of luck requires knowledge about the counter-factual choice under certainty. An experiment, consisting of a real-effort and a redistribution phase, can solve this causal inference problem associated with real-world data. In the different treatments, I vary the earnings procedure as well as the information provided about it while holding income associated with success (20 EUR) and failure (5 EUR) constant. That is, subjects are randomly matched into pairs such that one subject per pair experiences success while the other one experiences failure. In the redistribution phase, all participants can redistribute income from the subject that has received 20 EUR (success) to the subject that has received 5 EUR (failure) within another pair in the role of a spectator. In the EFFORT earnings condition, high and low income are assigned based on relative performance in one's pair. In the RANDOM earnings condition, high and low income are assigned randomly. In the FULL INFORMATION treatment, participants learn about the earnings procedure after the real-effort phase and before the redistribution phase. In contrast, they only know that both procedures are equally likely and identical for all subjects of the same session in the UNCERTAINTY treatment.

In the UNCERTAINTY treatment, high income spectators redistribute as much as in the RANDOM earnings condition with FULL INFORMATION whereas low income spectators redistribute as little as in the EFFORT earnings condition with FULL INFORMATION. Since the information provided and the resulting room for subjective beliefs is the only difference between the treatments FULL INFORMATION and UNCERTAINTY, I can interpret the effect of uncertainty on participants' beliefs directly from observing actual choices without having to rely on elicited beliefs. Consequently, the data strongly suggests that high (low) income spectators believe that luck (effort), rather than effort (luck), has determined everyone's income.

In the FULL INFORMATION treatment, high income spectators redistribute significantly less than low income spectators, irrespective of the earnings procedure (EFFORT or RANDOM). Moreover, all spectators redistribute more on average when earnings are determined randomly. Hence, the diametrically opposed effects of different experiences imply that distributional preferences of high and low income spectators are more aligned under uncertainty than under full information, though a significant difference exists in all treatments. Surprisingly, the majority of high income spectators do not equalize income but allow successful subjects to take home 50% more than unsuccessful subjects on average in the RANDOM earnings condition. Since success depends only on chance in this treatment, differences in distributive preferences can neither be due to different beliefs about the influence of luck nor due to self-selection. Rather, experience-dependent fairness ideals can explain the observed distributional preferences as a large fraction of 23% of high income spectators redistribute little or nothing to the poorer subject, in line with the libertarian fairness ideal. By contrast, low income spectators equalize income almost always. Survey responses given at the end of the experiment support this conclusion. Even when success is determined randomly and, thus, completely exogenous, the self-reported fairness ideals of high income subjects are more often libertarian than those of low income subjects.

Recently, Deffains et al. (2016) and Cassar and Klein (2017) concluded from their experiments that a self-serving bias drives differences in distributive preferences between high and low income earners. According to a self-serving bias, high income spectators attribute their success to internal factors such as hard work whereas low income spectators attribute their failure to external factors such as bad luck. With the present design, I can explicitly test for the presence of an effective self-serving bias. If high income spectators attributed their high income to their own effort while low income spectators attributed their low income to bad luck, the former should consider the effort-dependent earnings procedure more likely while the latter should tend to believe that earnings have been assigned randomly. Interestingly, I find exactly the opposite. That is, it strongly appears that a self-effacing bias, not a self-serving bias, affects beliefs and distributional preferences under uncertainty.

Since the difference in demand for redistribution between high and low income subjects is persistent across across treatments, a natural question to ask is whether only the experience of success, only the experience of failure, or the experience of both influence distributional preferences. In an additional information treatment, denoted by VOI, I observe participants' choices after they learn about the earnings procedure (EFFORT or RANDOM) but before they are informed about their own income.[2] First, I find that there is no difference in choices between those who (will) receive a high and those who (will) receive a low income, implying that the difference in preferences indeed comes from different experiences. Second, comparing the average transfer of all spectators from the VOI treatment with that of high and low income spectators in the differ-

[2]Since this environment resembles the concept of making a decision behind the veil of ignorance (Rawls, 1971), the treatment is abbreviated by VOI.

ent earnings procedures in the FULL INFORMATION treatments permits to observe what type of experience induces a change in distributional preferences. Indeed, I find that the experience of both success and failure changes distributional preferences. Depending on the earnings procedure, experiencing success reduces redistribution by high income spectators by 20 to 30% while it raises redistribution by low income spectators by 15 to 30%.

In the spirit of Smith (1759), economists, moral philosophers, and social scientists in general are often interested in the views of an impartial observer when evaluating the (in-)justice of unequal wealth distributions. Yet, as Konow (2008) notes, the concept of an ideal spectator by Smith (1759), who is fully informed and perfectly unbiased, does probably not exist in the real world. My findings strongly support his claim, showing that it holds even in clean and fully controlled laboratory environments. Specifically, the commonly used spectator method does not always elicit unbiased fairness views even though it eliminates the trade-off between self-interest and other-regarding preferences. Consequently, understanding how income affects fairness views both under certainty and uncertainty has important implications not only for public economics and political economy, but also for the design of experiments in general. More precisely, how should experimental economists deal with the effect of experience when they can control it? On the one hand, experience appears to be a prerequisite for informedness. How shall spectators assess the deservingness of high and low income, respectively, without experiencing the difficulty of a task themselves? On the other hand, the results from this experiment show that experience induces a bias towards the group the experience is being shared with. As the data from the VOI treatment suggest, the benefits from admitting experience can be obtained without inducing a bias by emplyong the spectator or third party observer method at the interim stage. That is, prior to learning the outcome but after participating in the real-effort phase, subjects are better informed and more likely to be unbiased.

3.2 Literature

This paper is most closely related to the experimental literature on people's fairness views. Employing an experiment that consists of an earnings and a redistribution phase, Cappelen et al. (2007) estimate that there is considerable heterogeneity in fairness ideals. While few participants act like libertarians who oppose any redistribution in their experiment, they conclude that the majority of participants either has a strict egalitarian or a liberal egalitarian fairness ideal. The fairness principle followed by liberal egalitarians is closely related to application of the accountability principle. Using a series of dictator games, Konow (2000) examines the extent to which the accountability principle can explain the willingness to redistribute and finds that whether dictators hold receivers accountable for their disadvantageous bargaining position explains a lot of the dictators' choices behavior. The experiment presented in this paper relies heavily on the prevalence of the liberal egalitarian fairness ideal. If uncertainty leads to different beliefs about the role of luck, it has an effect on people's choices only if support for redistribution indeed depends on those beliefs, as it should be the case for liberal egalitarians. Moreover, this paper also contributes to this literature by pointing out new explanations for the adoption of different fairness ideals.

Extending the above line of research, Cappelen et al. (2013) study fairness views with respect to risk taking by letting participants choose between risky and safe options in the first stage of their experiment. Subsequently, neutral spectators, who have not participated in the first stage, can redistribute income within pairs of participants with heterogeneous earnings. Generally, they find that redistribution is considerably lower if inequality prevails after participants chose differently risky options, even though most spectators compensate others for bad luck conditional on choices being identical. Mollerstrom et al. (2015) study fairness views under risk taking as well and find differences between environments involving controllable (chosen) and uncontrollable (forced) risk. Overall, these studies emphasize that preferences for redistribution depend on a number of different principles, some of which may be applicable only to quite specific earnings procedures. In contrast, the current paper focuses on the influence of experiencing either success or failure in general.

The effect of the earnings procedure is also explored in Durante et al. (2014) and Cappelen et al. (2010). In addition, the authors of the the latter work examine the effect of the participants' institutional background on their fairness considerations. In the sense that one's background and the chosen field of study is as assiciated with distinctive experiences in life, their work is related to mine. Their results suggest that business students with work experience are most likely to hold other people accountable for their economic situation. If the business school alumni indeed experienced more economic success than the other groups in their expexperiment, their findings would be in line with the conclusions drawn in this paper. Similarly, Almås et al. (2010) find that the strict egalitarian fairness ideal is more prevalent among children with a weaker economic background. However, because economic background or the choice of study field might be correlated with fairness ideals for others reason, these studies cannot identify any causal relationship.

Recently, the causal relationship between experience and distributional preferences has been experimentally investigated in two closely related studies. mIn Deffains et al.

(2016), subjects participate in a real-effort task to earn their endowments before making redistributive choices as spectators. They find that successful subjects accept significantly higher levels of inequality than unsuccessful subjects. They control for the difficulty of the task and, thus, for success, exogeneously and conclude that the difference in distributive preferences emerges from different beliefs about the importance of luck and effort in the task. Cassar and Klein (2017) also examine distributional preferences from spectators who have taken part in the earnings phase that leads to inequality between others. High or low income is assigned to subjects either randomly or based on relative performance. The chosen procedure varies across groups and is announced to all involved parties prior to the redistribution phase. Similarly to this paper, they find that distributive preferences depend on experience. Like Deffains et al. (2016), they conclude that either different beliefs about the role of luck or an in-group bias explains their results. However, they do not control for uncertainty exogeneously and, thus, do not test for an effective self-serving bias directly. Besides, they interpret their result in a way suggesting that only people experiencing failure but not those experiencing success are biased towards others who have experienced failure under the same procedure. My conclusion based on the comparison of the treatments FULL INFORMATION and VOI is different. In contrast, I find that distributional preferences of both successful and unsuccessful subjects are strongly influenced by their experience.

Though seemingly similar, the self-serving bias is different from the self-confirmation bias invoked by Konow (2000). The latter results from the desire to reduce cognitive dissonance when facing a trade-off between choosing a fair allocation and maximizing own material benefits. In contrast, a self-serving bias emerges when individuals attribute failure to external factors beyond their control but attribute success to their own effort and choices (Miller et al. (1975), Bradley (1978)).

Besides, there is a large empirical literature documenting the relationship between an individual's belief about the role of luck for economic success and her or his support for redistribution. Based on survey data from 12 countries, Corneo and Grüner (2002) find that not only income, but also a respondent's belief that *"hard work is key"* negatively affects support for redistribution. Fong (2001) presents similar findings from another survey. Likewise, Alesina and La Ferrara (2005) show that both the expected future income and perceived equality of opportunity explain support for redistribution. The present work contributes to this literature by exploring both a potential source of those heterogeneous beliefs as well as their causal impact on distributional preferences.

A positive relationship between the degree of luck involved in the earnings process and the support for redistribution is also at the core of several theoretical works on the political economy of redistribution. Piketty (1995) models household income as a process that depends on luck, effort, and ancestry. In equilibrium, dynasties with bad luck develop more pessimistic beliefs, thus, exert less effort and, still, strongly support redistributive politics. Alesina and Angeletos (2005) show that a society may either be in an equilibrium with a high dependence of income on luck, high redistribution and low effort, or in an equilibrium with a limited role of luck, little redistribution and high effort. Although the experiment is not designed to explicitly test any of these theories, it sheds light on a mechanism underlying all of them, the formation of beliefs about the role of luck under uncertainty.

Finally, this paper contributes to the debate about the the relationship between income and giving. From an empirical perspective, the effect of income on the willingness to give seems ambiguous. Some studies find a positive relationship (Eckel et al., 2007), some find a U-shaped relationship (Auten et al., 2000), and others find no relationship between income and giving (Andreoni and Vesterlund (2001) or Buckley and Croson (2006)). This opposed evidence might partially be due to confounding factors influencing the decision to give. One of these factors has recently been identified by Erkal et al. (2011) who report evidence from a real-effort experiment in which more pro-social agents exert less effort than less pro-social agents. The relatively modest willingness to give among the richest subjects they observe in the data thus is the result of a self-selection of less pro-social subjects into high income ranks. The potential dependence of distributional preferences on experience, which is the focus of this paper, constitutes another possible explanation.

3.3 Experimental Design

The experiment has two main stages. The first stage is a production phase involving a real-effort task, which is identical across all treatments. Whether income is assigned to subjects based on their effort or randomly is determined only after the first stage is completed. In the second stage, participants can redistribute earnings between two other participants as spectators. Finally, all participants answer several belief and survey questions.

3.3.1 Real-Effort Task

At the beginning of the real-effort task, participants are matched into pairs. Depending on the treatment, which is not revealed to participants before the real-effort task is completed, relative performance within a pair determines the pre-redistribution stage earnings. The task requires participants to encrypt fictional and randomly generated words by using a code key which assigns a three digit number to each letter of the latin alphabet. Both the word that has to be encrypted and the code key change once a word is encrypted correctly. The change of the code key not only concerns the numbers that have to be used to translate each letter, but also the ordering of the letters. The task has been developed by Benndorf et al. (2014) to minimize learning and is essentially a modified version of the encryption task used in Erkal et al. (2011). During the task, participants can keep track of the number of correctly encrypted words. If the entered solution is incorrect, the coding table does not change and they have to try again. Performance in the task yields a score, which is measured as the number of correctly encrypted words. Figure (B.1) in the appendix shows a screenshot of the task. All participants have 25 minutes in total to work on the encryption task.

3.3.2 Redistribution Task

Upon completion of the real-effort task, one subject in a pair receives 20 EUR while the other one receives 5 EUR. In the redistribution phase, those earnings are subject to change depending on other participants' redistributive choices. This is because every subject is given the opportunity to change the distribution of earnings within another pair. That is, they can choose to transfer any amount between 0 EUR and 15 EUR (in 10 cents increments) from the high income participant to the low income participant of another pair. This range is chosen such that each participant is secured a minimum pay of 5 EUR.[3] To select their preferred transfer, subjects can use the slider displayed at the bottom of their screens (see figure (B.2) in the appendix). For every interim position of the slider, the screen visualizes how the transfer implied by the slider's position affects the other subjects' earnings. Participants have to confirm the final slider position before moving on to the belief elicitation task. When making their redistribution decision, subjects only know their own score from the encryption task but never those from the subjects affected by their redistribution decision.

The instructions (see section B.3) emphasize to participants that their decisions are made for another pair and do never affect themselves. They also know that it is impossible that their own earnings are affected by the decision of someone from the group for which they choose a transfer. In addition, it is also explained to them that their choice is implemented only with a probability of 50%. This is because when all

[3] Negative transfers that shift earnings from the poorer to the richer subject are not allowed to ensure a minimum pay to all participants without requiring a show-up fee. Yet, exactly the same outcome space would obtain even without restricting transfers by paying a show-up fee of 5 EUR to all participants and by setting the high income equal to 15 EUR and the low income equal to 0 EUR.

subjects make a choice for exactly one pair, there are twice as many decisions as there are pairs.

3.3.3 Treatments

The treatments variations in this experiment are organized along two dimensions. The first dimension concerns the conditions that determine who receives the high income of 20 EUR and who receives the low income of 5 EUR in pair. In the effort condition (E), the subject with the higher score receives 20 EUR while the subject with the lower score receives 5 EUR.[4] In the random earnings condition (R), high and low incomes are assigned randomly within each pair. At the beginning of the experiment, participants know only that both earnings procedures are equally likely and that all participants of the same session are subject to the same procedure. The second dimension concerns information that can be of three different kinds. In the full information treatment (FI), subjects learn about the earnings procedure that is chosen for their session before they make their redistributive choices. Also, they learn about their own income from the real-effort phase when making the redistribtuion decision. In the second information treatment, subjects receive no information about how income has been awarded and thus remain uncertain about the underlying earnings procedure (UP). Yet, they still know about their own income from the real-effort phase while choosing the transfer for the other pair. In a third treatment, subjects learn about the earnings procedure but receive no information about their own income prior to making a decision in the redistribution task. As subjects in this treatment decide without knowing where they stand, it resembles the veil of ignorance condition suggested by Rawls (1971) and is therefore abbreviated by VOI.

To summarize, the experiment employs a 2×3 design. The earnings conditions effort (E) and random (R) are combined with one of the three information treatments, which are full information (FI), uncertain earnings procedure (UP) and veil of ignorance (VOI). The information treatments FI and UP are the main treatments in this experiment as the comparison of them permits to investigate the effects of uncertainty in the light of heterogeneous experiences of economics success. Because subjects might be biased in favor of participants who experience the same outcome as they do, the VOI treatment offers an interesting benchmark that allows to observe distributional preferences of a spectator who cannot associate her or himself with either subject in the other pair.

[4] Recall that a subject's score is his or her number of correctly encoded words.

3.3.4 Belief and Survey Questions

Immediately after completing the redistribution task, participants are informed about the outcome of the redistribution phase. For participants in the VOI treatment, this means that only then, they learn about both their income from the real-effort phase and the transfer chosen for their pair. Only afterwards, subjects' beliefs about the impact of luck in the earnings procedure are elicited. Upon completion of the belief elicitation stage, subjects respond to two additional survey questions regarding their general attitudes towards inequality and redistribution.

Belief elicitation. In all treatments, subjects indicate how strongly they believe that the high income of 20 EUR has been allocated to the subject with the higher score within their own pair. That is, a participant who earned 5 EUR before the redistribution phase is asked whether he believes that his score is below his opponent's whereas a participant who earned 20 EUR is asked whether she believes that her score is above her opponent's. In addition, subjects are asked a second belief question if they are in the UP treatment, where they have not received information about whether effort or luck has determined income. The question asks subjects how certain they are that income has been based in relative performance on their session or, in other words, whether they believe that they are in treatment E. For each question, subjects could choose a probability between zero and one hundred percent in increments of ten percent. Figure (B.3) and figure (B.4) in the appendix contain the exact questions as displayed on the subjects' screens.

The belief elicitation is not incentivized. Naturally, this increases the risk of not measuring true beliefs, for example because stating particular beliefs might allow the respondent to improve his or her self-image. However, incentivizing respondents by rewarding them if their stated belief is sufficiently correct may induce participants to reevaluate the situation. Yet, I am interested in what participants (subconsciously) believe while making their decision. Since additional incentives are likely to increase the gap between actual beliefs when making a decision and stated beliefs in this experiment, they are omitted.

Survey. The first question addresses participants' fairness ideals. The first part of the question describes a situation in which there is great economic inequality between two individuals who are about to retire. The second part asks under what conditions participants would be against redistribution between the two, offering four different choices, excluding "no opinion"(see figure B.5). The available options range from (I) unconditional opposition over opposing redistribution only if either (II) mostly effort and individual choices have determined income or if (III) both individuals had equal opportunities and mostly effort and individual choices have determined income to (IV) unconditional support for redistribution. Roughly speaking, each of these options corresponds to a particular fairness ideal. Notably, I offer subjects two characterizations of the liberal egalitarian fairness ideal, one that includes equality of opportunity as a necessary condition for no redistribution and one that does not, in order to make the classification less coarse. Additionally, subjects indicate how much they believe economic success in Germany is driven by luck versus hard work in the second survey question. Similar questions exist in the General Social Survey (GSS) or the International Social Survey Programme (ISSP), allowing me to observe how the student sample under consideration here compares with a more representative sample of the population.

Asking participants about their fairness views and beliefs about the role of luck in the real world also serves another purpose. If there are no treatment effects, it could be because all participants are strict libertarians who oppose any form of redistribution, irrespective of their experience or of the preceding earnings procedure. Similarly, people who are convinced that everything depends on luck are likely to always support redistribution, again inhibiting any treatment effects. Responses may thus by used to explain how likely a null result is due to the composition of the subject pool.

3.3.5 Theory: Rational and Biased Beliefs

In the full information treatment, subjects know with certainty whether income in their session has been allocated randomly (R) or based on relative performance (E). If the average transfer of subjects experiencing success in treatment UP matches the transfers of successful subjects in treatment FI for a particular earnings procedure, say E, then the belief of those subjects experiencing success in treatment UP must be similar to that of successful subjects in treatment FI + E. Since the latter group knows that earnings are determined by relative performance with certainty, I would be able to infer that the beliefs of high income subjects put a high weight on earnings procedure E. Consequently, analyzing *what* beliefs about the role of luck participants form under uncertainty requires only that demand for redistribution be monotonous in the belief about the role of luck, which is a common finding in the empirical literature described in section 3.2. However, inferring the existence of particular biases from those beliefs requires to know what a Bayesian decision maker would believe. Importantly, the Bayesian posterior belief in treatment UP is not necessarily identical to the prior since the combination of both a subject's income from the encryption task and her score constitute an imperfect signal about the true earnings procedure.

The main predictions concerning beliefs of a Bayesian individual about the earnings procedure in UP are: Conditional on receiving the high income, E is more likely than R if and only if the achieved score lies above the median score.[5] Analogously, R is more likely than E if and only if the achieved score lies above the median score conditional on receiving a low income. A formal proof of these statements is presented in section B.1, though both results are quite intuitive. For example, consider a subject who knows that her score is below the median score in the population. Then, her probability of receiving the high income is below 50% under condition E but equal to 50% under condition R. Hence, because receiving the high income is more likely under condition R, the posterior attaches more weight on R when she receives the high income. The behavioral implications for the present experiment are summarized below.

Observation 1 *If subjects form rational beliefs, distributional preferences of those receiving 20 EUR in UP should be more similar to those in FI + E if their obtained score is above the median score and be more similar to FI + R otherwise.*

Observation 2 *If subjects form rational beliefs, distributional preferences of those receiving 5 EUR in UP should be more similar to those in FI + R if their obtained score is above the median score and be more similar to FI + E otherwise.*

Importantly, I do not expect subjects to display behavior that is in line with Bayesian updating. Instead, there are two potential biases which might be present in such an environment according to the existing literature in psychology and experimental economics. First, actual beliefs and distributional preferences might deviate from the Bayesian benchmark due to a self-serving bias. If subjects are prone to a self-serving bias, they tend to attribute success to internal factors such as own performance. Conversely, they attribute failure to external factors such as bad luck. Second, participants might exhibit a self-effacing bias, which is characterized by the opposite behavior and thus leads to opposite predictions. Prediction 1 and 2 summarize the main implication of both biases regarding distributional choices in treatment UP.

Prediction 1 *If a (strong) self-serving bias influences distributional preferences, subjects with a score below the median score display the following behavior: low income subjects choose transfers similar to the transfers chosen by low income subjects in treatment FI + R and different from those in treatment FI + E. Conversely, high income subjects choose transfers similar to the transfers chosen by high income subjects in treatment FI + E and different from those in treatment FI + R.*

Prediction 2 *If a (strong) self-effacing bias influences distributional preferences, subjects with a score above the median score display the following behavior: low in-*

[5] It is assumed that a Bayesian individual knows the distribution of scores in the population

Information Treatment	FI	FI	UP	UP	VOI	VOI
Earnings condition	E	R	E	R	E	R
Observations	64	60	32	26	16	12

Table 3.1: Number of participants per treatment

come subjects choose transfers similar to the transfers chosen by low income subjects in treatment FI + E and different from those in treatment FI + R. Conversely, high income subjects choose transfers similar to the transfers chosen by high income subjects in treatment FI + R and different from those in treatment FI + E.

Another related behavioral bias known as overconfidence would have implications similar to those of a self-serving bias. As overconfidence is associated with a too optimistic belief concerning own performance, it can be formally modeled as a prior that understates the median score in the population. As a consequence, high income subjects would interpret even scores below the median as signals of condition R and vice versa. Conversely, the effect of underconfidence is qualitatively equivalent to that of a self-effacing bias.

3.3.6 Subject Pool and Experimental Procedures

In total, 210 subjects participated during 14 sessions in the experiment from June 27 to July 12, 2017. Subjects were recruited via hroot, between 12 to 18 subjects participated in each session. Each session lasted about 50 minutes and average payoffs amounted to 12.50 EUR. Payoffs were composed only of earnings and chosen transfers, there was no show-up fee. Participants were students from the University of Heidelberg and came from a wide range of study fields, with one quarter having a background in economics. The language of the experiment was German.

Table 3.1 reports the number of subjects who participated in each of the treatment conditions. Different numbers for different earnings conditions within the same information treatment are due failure to show up of registered participants. The differences between information treatments are on purpose however. In treatment UP, subjects do not know whether they are in condition E or R. Consequently, their choices can be considered jointly, making this group of participants just as large as the groups their choices are to be compared with: FI + E and FI + R. Similarly, the choices of high and low income subjects are comparable in the VOI information treatments, allowing me to merge the data as well.[6]

3.4 Results

I report the results from the real-effort task first. Next, results on distributional preferences are presented in the following order: redistribution under full information (FI), redistribution behind the veil of ignorance (VOI), and redistribution under uncertainty (UP). Reported p-values are based on the non-parametric two-sided two sample Mann-Whitney-Wilcoxon (or Mann-Whitney U) test.

3.4.1 Real-Effort Task

The average number of encrypted words (score) from all treatments is 109 with scores ranging from 5 to 167 words and the median score is 108 words. Time-dependent data show that very low scores result from participants stopping to work on the encryption task early, not from being unable to encrypt words in general. In total, this happened three times. Figure (3.1a) shows the average score achieved across treatments. Note also that the apparent oversampling of high score participants in treatment condition FI + E does not affect the interpretation of results as the subsequent analysis shows that own performance has no effect distributional choices. According to figure (3.1b), participants were equally productive in all rounds. That is, there was no general tendency for participants giving up too early nor becoming better at the task over time. This implies that learning was successfully prevented, thus making it truly costly for subjects to continue providing effort until the end.

3.4.2 Redistribution Under Certainty

In this section, I present the results for the effort and random treatment condition under full information. Henceforth, I use the notion of high or low income spectator rather than high or low income subject to distinguish between the subject making a decision and the subject being affected by it. Recall that in both treatments, decision makers are spectators who have no stake in the decision, implying that their chosen level of redistribution should depend only what they consider a fair distribution. In addition, the earnings procedure is identical for all participants of the same session. If

[6] The results presented in the next section confirm that there are no differences in both cases.

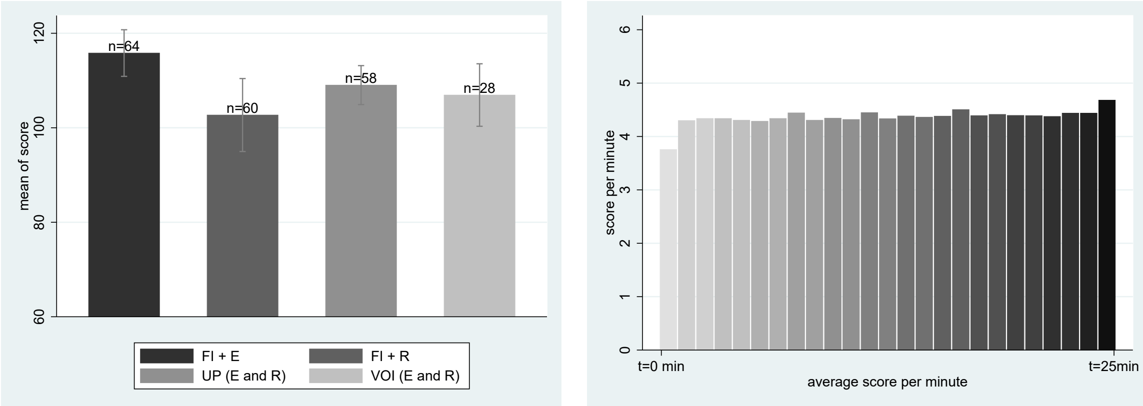

(a) Score by treatment (b) Effort over time

Confidence intervals in figure (a) are at the 90% level.

Figure 3.1: Performance encryption task

fairness views about inequality depend only on the process leading to it, there should be no differences in distributive preferences between high and low income spectators.

However, average transfers to low income subjects vary greatly and systematically with the spectator's income. For the treatment FI + E, where income depends on effort, figure (3.2) shows that high income spectators redistribute 3.58 EUR on average. In contrast, low income spectators transfer 6.71 EUR on average, implying a relative difference of almost 90%. Possible explanations for why there are such large and significant differences in the preferences of spectators are discussed in great detail in section 3.5. Briefly, an important factor is that the view held about whether the procedure determining income is fair depends on one's own outcome. That is, even though income depends on effort in treatment FI + E, high income spectators were still lucky since the choice of the encryption task was in their favor. Because the choice of a particular task always creates unequal opportunities, treatment FI + E still leaves room for heterogeneous judgments concerning the influence of luck.[7] Yet, inequality of opportunity does not prevail in treatment FI + R as all participants have an equal chance of "winning" the high income. While this does certainly not imply that people should have no demand for redistribution in R, it means that there is no room for reasonably different views about the role of luck across high and low income spectators. Hence, it seems implausible that distributional preferences depend on own success or failure when success is knowingly the outcome of a purely random process. Yet, a significant difference in the chosen transfers persists between high and low income spectators as displayed on the right-hand side of figure (3.2). Low income spectators redistribute 7.40 EUR on average and often choose a transfer of 7.50

[7]Nevertheless, the effect of luck is of course strictly smaller than the random earnings condition (R).

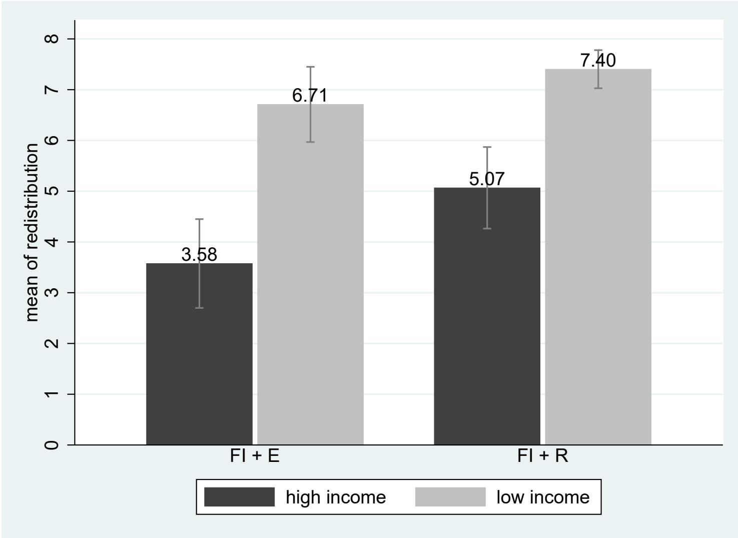

Differences between high/ low income are significant at $p < 0.0001$ both in FI + E and in FI + R. Differences between E and R are significant at $p < 0.01$ both conditional on high and low income. Note that a transfer of 7.50 EUR equalizes earnings. Confidence intervals are at the 90% level.

Figure 3.2: Average transfers under full information

EUR, thereby completely equalizing both subjects' earnings. In contrast, high income spectators choose an average transfer of 5.07 EUR, which is about 30% lower. Looking at the distribution of individual choices reveals that only 12 out of 30 high income spectators choose a transfer that (almost) equalizes earnings as opposed to 25 out of 30 low income spectators.[8]

Recall that in the zero merit environment that treatment FI + R refers to, the incomes of both the spectator and the subjects from the pair affected by the spectator's decision are assigned randomly. Consequently, the high income of the spectator is completely exogenous and there are no systematic differences between high and low income subjects a spectator could rely on to discriminate. Thus, while results from the existing literature suggest that people follow a luck-egalitarian fairness ideal and do not accept any inequality in environments that are purely luck-dependent, the opposed evidence from this experiment suggests that experience has a fundamental influence on whether that is true or not.

Regarding the comparison of conditions E and R, I find that redistribution is generally higher when income is assigned randomly. Compared to FI + E, high income spectators redistribute 1.49 EUR more on average in FI + R. For low income spectators, the difference amounts to 0.69 EUR only, because redistribution is already high in FI + E. The observed difference is significant in both cases ($p < 0.01$). Notice that existing research on fairness views has mostly focused on the difference in demand for

[8]The complete distribution of transfers is displayed in figure (3.5).

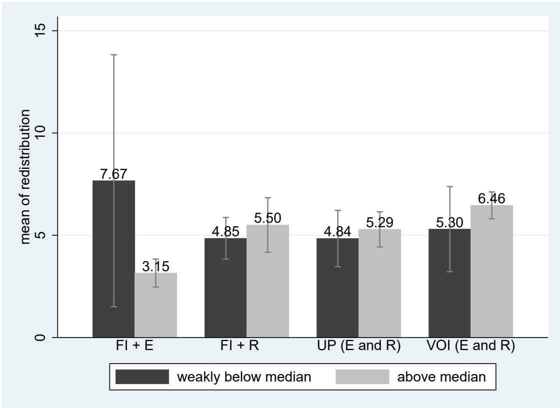

(a) High income

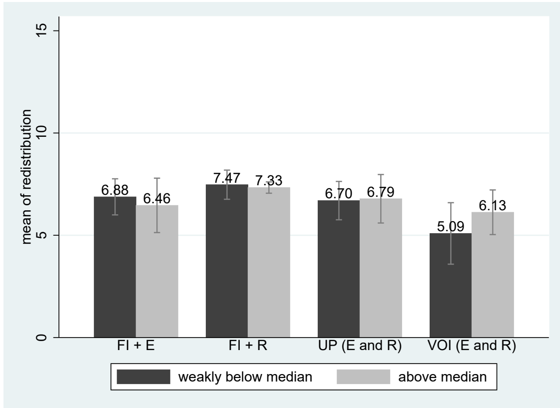

(b) Low income

Figure (a) displays average redistribution from high income spectators, (b) from low income spectators. Confidence intervals at the 90% level. There are only three high income spectators with below median income in FI + E, preventing the construction of a reasonable confidence interval.

Figure 3.3: The effect of score on redistribution across treatments

redistribution induced by different earnings procedures. In this experiment, the average spectator redistributes 1.09 EUR more in treatment R than in treatment E (1.49 EUR for high and 0.69 EUR for low income spectators). In comparison, the average difference between the transfers of high and low income spectators amounts to 2.33 EUR in treatment R and to 3.13 EUR in treatment E. Even though quantifying effect sizes of course depends on the experimental design, those results strongly suggest that the effect of experiencing economic success on distributional preferences can be way larger than that of the underlying earnings procedure.

Endogeneity in treatment condition E.

Since high and low income are completely exogenous in treatment FI + R, the effect of experiencing economic success on distributional preferences has to be be causal. In all treatments relying on the effort based earnings procedure (E) however, any correlation between experience in distributional preferences need not be causal per se. Instead, those participants who oppose redistribution in general could simply be types that are more competitive, making them more successful in the encryption task.

Figure (3.3) contains evidence against such a claim. If subjects' performance was correlated with their general attitude towards redistribution, higher scores should correlate with lower transfers even after conditioning on the participant's income. Yet, across all treatments, transfers by spectators who score above the median are about the same or even higher than those by spectators with a score below the median (see figure 3.3a). Similarly, figure (3.3b) shows that redistribution does not depend the on score for low income spectators either.

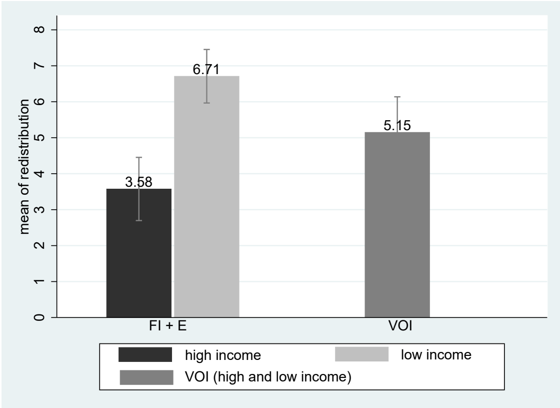

(a) Earnings procedure E

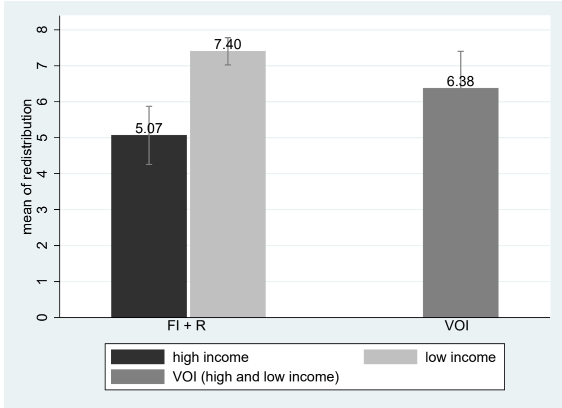

(b) Earnings procedure R

In figure (a), the difference between high income (FI + E) and VOI is significant ($p < 0.01$). A comparison between veil of ignorance (voi) and low income suggests a difference ($p = 0.11$). In (b), the difference between low income (FI + R) and VOI is significant $p = 0.02$). The comparison between VOI and high income suggests a difference ($p = 0.13$). Confidence intervals at the 90% level.

Figure 3.4: Redistribution in UP and VOI

Moreover, the results from the VOI treatment show that there is no difference in transfers between high and low income spectators when spectators do not know their own income. Thus, the differences between high and low income spectators observed in treatment E do not seem to be due to self-selection but due to the experience of success or failure.

3.4.3 Redistribution Behind the Veil of Ignorance

The treatments VOI + E and VOI + R permit to observe distributional preferences before they may be influenced through the experience of success or failure. Since the results displayed in figure (3.3) show that transfers are independent of scores and no information other than their own score is disclosed to spectators in the VOI treatments, I analyze the data from high and low income spectators jointly. That is, the single right-hand side bars in figure (3.4) contain transfer choices by all spectators in treatment VOI + R and VOI + E, respectively. According to figure (3.4a), the transfer on average equals the mean of transfers chosen by both high and low income spectators in treatment FI + E, implying that both groups' preferences are influenced through their experience to the same extent but in opposite directions.

As figure (3.4b) shows, the same pattern can be observed in treatment VOI + R. Again, the average of all transfers in FI + R compares with the mean transfer chosen behind the veil of ignorance under the random earnings procedure. In contrast to the behavior of high income spectators in FI + R, luck-egalitarianism becomes the

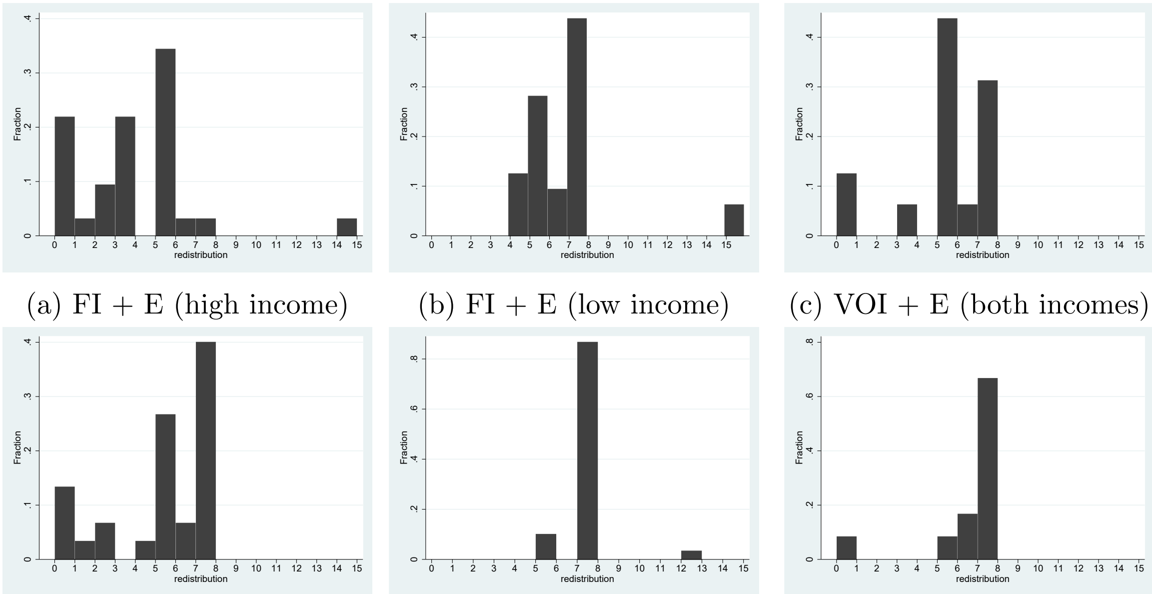

(a) FI + E (high income) (b) FI + E (low income) (c) VOI + E (both incomes)

(d) FI + R (high income) (e) FI + R (low income) (f) VOI + R (both incomes)

Histograms of individual transfers with data organized such that the bar between x Euro and x+1 Euro contains all observations with transfers t satisfying $t \in [x, x+1)$.

Figure 3.5: Relative frequencies of transfers

predominant fairness ideal in the absence of any experience effects as figure (3.5f) shows. That is, 6 out of 12 spectators choose a transfer of 7.5 EUR to equalize earnings. Another two spectators choose 7 EUR, implying that roughly two thirds of spectators implement almost perfect equality when their own outcome is not yet revealed to them. Interestingly, the distribution depicted in figure (3.5f) compares considerably well with the distribution of transfers by low income spectators in figure (3.5d), again suggesting that high income spectators are affected by their experience more often than low income spectators in treatment FI + R.

As mentioned in the previous section, the data shown in figure (3.4a) also addresses the endogeneity question concerning the findings from treatment FI + E. Since the timing of the income information constitutes the only difference between FI + E and VOI + E, the observed differences in redistribution levels must be driven by experience of the outcome rather than by self-selection.

3.4.4 Redistribution Under Uncertainty

I here present data from the UP treatments in which subjects remain uncertain about the earnings procedure. As explained in section 3.3.4, comparing the transfers under uncertainty with those chosen under full information allows to make an inference about

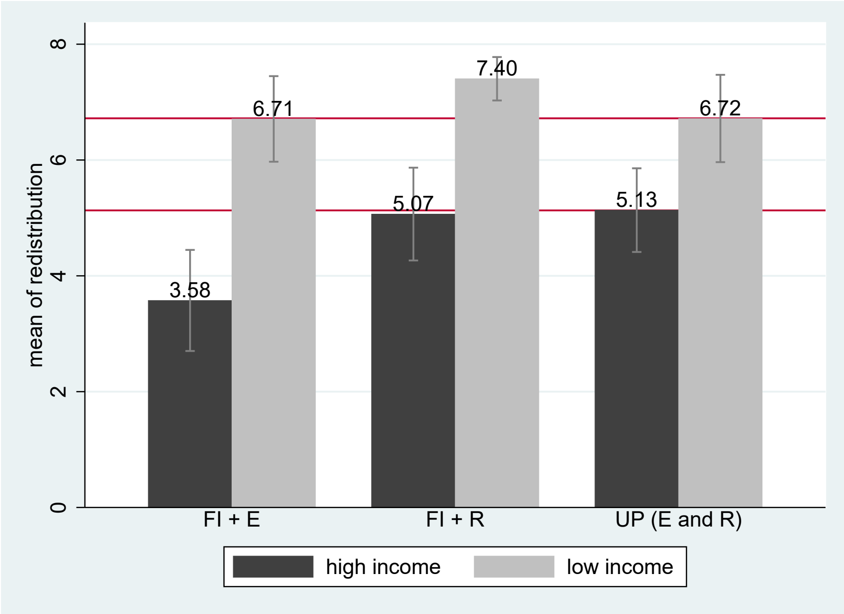

The difference between high income (UP) and high income (FI + E) is significant at $p < 0.01$. The difference between low income (UP) and low income (FI + R) is also significant at $p < 0.05$. Confidence intervals are at the 90% level.

Figure 3.6: Average transfers under certainty

subjects' beliefs. For the first part of this analysis, I analyze the data from the different earnings conditions E and R jointly.[9]

Figure (3.6) contains a striking result. Under uncertainty, high income spectators choose transfers as in the full information environment with a randomness-based earnings procedure (FI + R). In other words, the earnings distribution considered as fair by high income spectators under uncertainty equals the earnings distribution considered as fair by high income spectators who know that income has been determined based on pure chance. Conversely, the average amount redistributed by low income spectators under uncertainty almost equals the amount transferred by low income spectators who know that earnings are based on effort (FI + E). Notably, the difference with respect to the other treatment (FI + R for low income spectators and FI + E for high income spectators) is always significant. That is, high income spectators distribute significantly more to the poorer subject under uncertainty ($p < 0.01$) than when they know for sure that the earnings procedure is based on effort (FI + E). Besides, low income spectators redistribute significantly less to the poorer subject ($p < 0.05$) than when they know for sure that luck has determined income (FI + R), implying that the

[9]The data shows that transfers are independent of the actual earnings procedure both for high and low income spectators. This finding seems in line with evidence presented in the previous section. Even though the fraction of spectators with an above median score among high income spectators is larger in treatment UP + E than in treatment UP + R, there is no significant difference because transfers are completely independent of scores.

difference between high and low income spectators under uncertainty is significantly smaller than in any of the full information treatments.

At first glance, this behavior seems incompatible with the notion of a self-serving bias. Instead of inferring from their own success that relative performance must have determined income, high income spectators seem to act exactly in the opposite way. Yet, interpreting the above results in favor of (or against) a particular bias would be premature. This is because in order to formally reject the influence of a self-serving bias, I must reject that the average transfer of high income spectators in UP is larger than the transfer a Bayesian high income spectator with an average score would choose (and vice versa for low income spectators). Predicting what the average transfer should be if spectators formed Bayesian beliefs about the earnings procedure depends both on the exact distribution of achieved scores among high and low income spectators and on how beliefs map into transfers. Since the distribution of scores is random and the mapping from beliefs to transfers unknown, I use what is probably the best available heuristic: I compute the (equally weighted) average of transfers chosen by fully informed high (low) income spectators in both treatments FI + R and FI + E.

For high income spectators, the heuristically computed average amounts to 4.33 EUR. Analogously, I obtain an average transfer of 7.06 EUR for low income spectators. By constructing a hypothetical sample in which all high income spectators' choices equal 4.33 EUR and all low income spectators' choices equal 7.06 EUR, I test whether the transfers chosen under uncertainty differ significantly from the predicted Bayesian benchmark (BB). The non-parametric test shows that the average transfer of high income spectators under uncertainty (5.13 EUR) is still significantly larger than the benchmark of 4.33 EUR. This result strongly rejects a self-serving bias and provides evidence in support of a self-effacing bias among high income spectators. For low income spectators, the difference between the average transfer under uncertainty (6.72 EUR) and the computed benchmark of 7.06 EUR is significant as well, implying that low income spectators also display no self-serving but a self-effacing bias. For reference, all p-values computed for the above comparisons are displayed in table 3.2.

The relative frequencies of transfers across treatments as displayed in figure (3.7) support the general picture of high income spectators under uncertainty behaving as in treatment FI + R and low income spectators behaving as in FI + E. In contrast to the apparent differences between high and low income spectators in the FI treatments (see figure 3.5), the distributions of transfers by high and low income spectators under uncertainty are strikingly similar. In line with the comparison of the average transfers, this shows that uncertainty harmonizes distributive preferences, making support for redistribution across different income classes more aligned. Moreover, the congruence

Income	low	low	low	high	high	high
UP compared with	FI+E	FI+R	BB	FI+E	FI+R	BB
p values	0.4243	0.0313	0.01094	0.0045	0.6577	0.0000

The first row indicates whether the test includes observations from low or high income spectators. The second row indicates the treatment combination that the choices of low/ high income spectators from treatment UP are compared with. Reported p-values are based on a two-sided, two-sample Mann-Whitney U test.

Table 3.2: Comparison of FI and UP treatments

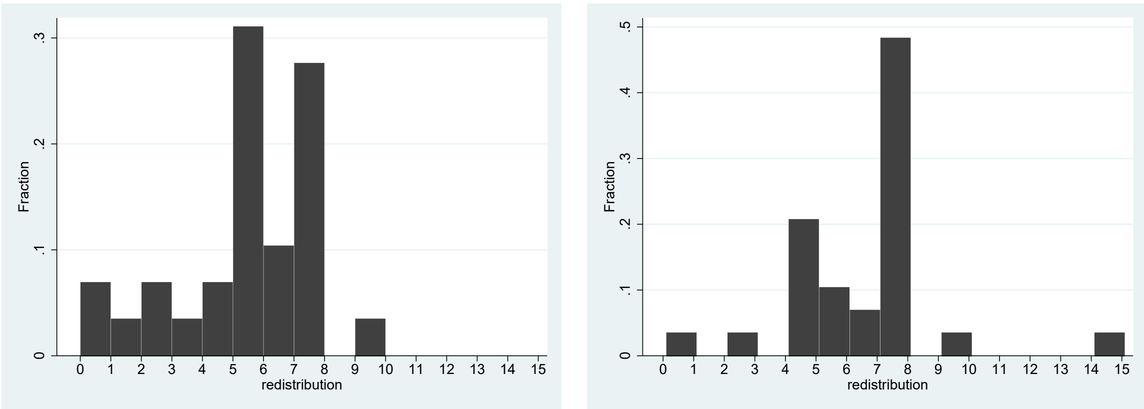

(a) High income (b) Low income

Histograms of individual transfers with data organized such that the bar between x Euro and x+1 Euro contains all observations with transfers t satisfying $t \in [x, x+1)$.

Figure 3.7: Distribution of individual transfers under uncertainty (UP)

of distributions displayed in figure (3.7a) and (3.5d) as well as in figure (3.7b) and (3.5b) is remarkable, thus confirming the above conclusion about the influence of a self-effacing bias. To summarize the overall comparison between UP and FI, I find that uncertainty harmonizes distributive preferences, making support for redistribution across different income classes more aligned.

Additional evidence against a self-serving bias can be found in the data by contrasting the results with prediction 1 and 2. Figure (3.8) shows the average transfer of participants divided by the median score across the treatment combinations FI + E, FI + R, UP (both R + E) conditional on either (a) high or (b) low income.[10] According to prediction 1, high income spectators with a below median score should (still) believe that their success is the result of high performance, making them consider condition E

[10]Observations from the VOI treatment are not included here because transfers cannot reasonably be conditioned on income when subjects have no information about their income.

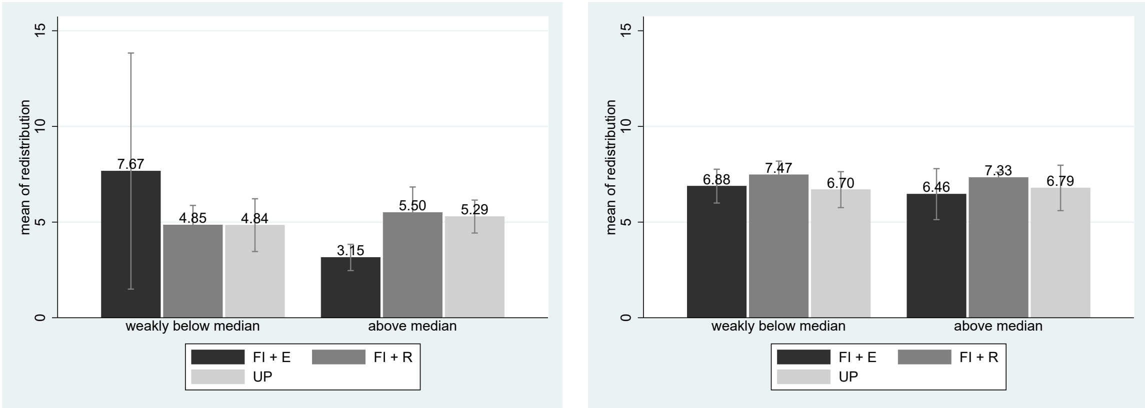

(a) High income (b) Low income

Figure (a) shows the average transfer of high income spectators divided by score (above or below median) and by treatment. Figure (b) shows the analogue information for low income spectators. Confidence intervals at the 90% level.

Figure 3.8: Redistribution by score and treatment

more likely than R. As a result, their transfers should be more similar to those of high income spectators in treatment FI + E. Figure (3.8a) shows that conditional on a below median score, the average transfer of high income spectators in treatment UP perfectly matches that of high income spectators in treatment FI + R (4.85 EUR and 4.84 EUR). Analogously, transfers chosen by low income spectators with a below median score are very similar to the transfers of low income spectators with a below median score in treatment FI + E (6.70 EUR and 6.88 EUR) as depicted in figure (3.8b). Both findings stand in sharp contrast to prediction 1 and, thus, to the implications of a self-serving bias.

Yet again, the data fully supports prediction 2 and the presence of a self-effacing bias. Conditional on a score above the median score, I observe that redistributions by high income spectators in UP are similar to redistributions by high income spectators in treatment FI + R (5.29 EUR and 5.50 EUR, see figure (3.8a)). Likewise, low income spectators with an above median score redistribute 6.79 EUR on average, which is closer to the average of 6.46 EUR of low income spectators in FI + E than to the average of 7.33 EUR of low income spectators in FI + R (see figure 3.8b).

3.4.5 Elicited Beliefs

Beliefs in the belief elicitation task are measured on a nominal scale using percentages. Recall that the question that is asked across all treatments concerns the probability people attach to the event that relative income in their own pair reflects relative per-

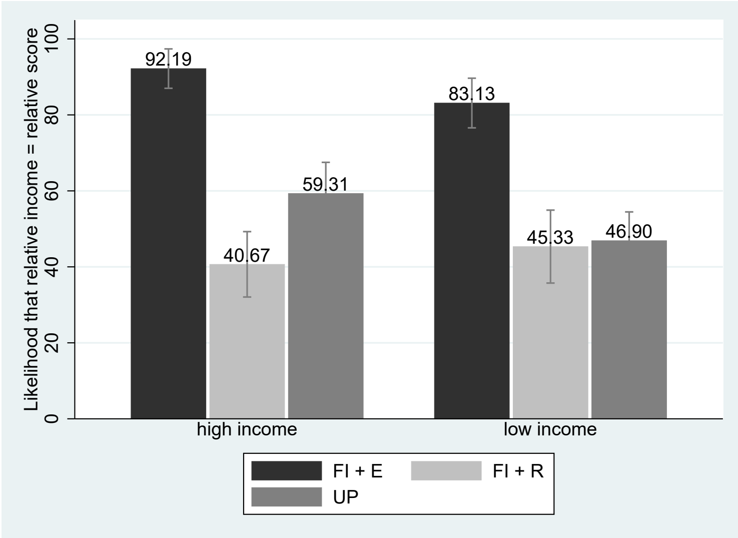

Question: "What do you think is the probability that your earned income coincides with your relative performance in the tournament?" Figures are in percentages. Confidence intervals are at the 90% level.

Figure 3.9: Elicited beliefs about the influence of luck

formance.[11] That is, if a subject is 100% certain that her performance was higher in her pair while she has received 5 EUR only, she chooses 0%. Alternatively, a subject that is undecided about that question may select 50% and so forth.

Subjects' beliefs elicited through this first question are displayed in figure (3.9). In treatment FI + E, both high and low income subjects should be certain that their income has been the result of relative performance. Stated beliefs are quite close to 100% and much higher than in the other treatments, verifying most participants trusted the instructions and understood the earnings conditions well. In treatment FI + R, the distribution of income is completely random. Thus, relative performance coincides with relative income in only about 50% of all pairs, which should be the average belief if beliefs were unbiased. The corresponding bar shows a magnitude of slightly less than 50%, indicating that subjects might have been slightly underconfident.

Lastly, observe that beliefs of high income subjects under uncertainty (UP) roughly match the average of all high income spectators' beliefs in FI + I and FI + R, which appears reasonable without any further assumptions. However, the average belief of low income subjects in treatment UP is just as low as in FI + R. At first glance, this may seem to be in line with a self-serving bias, which Deffains et al. (2016) infer from elicited beliefs as well. However, the data depicted in figure (3.10a), which again

[11] The ensure participants understood this question, the following explanation was displayed below the question: "If you have earned 5 EUR, we would like to know how certain you are that you had the lower score in your group. If you have earned 20 EUR, we would like to know how certain you are that you had achieved the higher score in your group?".

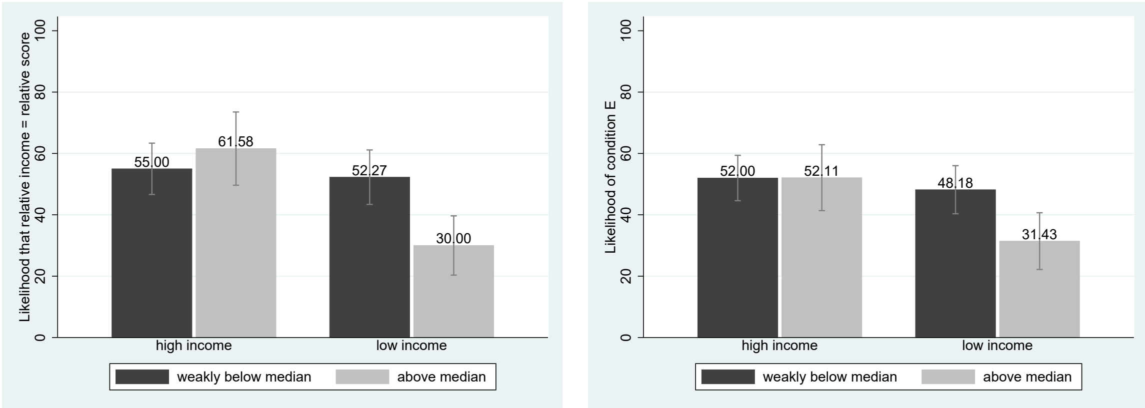

(a) Belief about own outcome (b) Belief about session (E or R?)

Figure (a) reports elicited beliefs about the influence of luck as explained in figure (3.9). Figure (b) reports elicited beliefs from the following question: "What to you think is the likelihood that your session was part of world 1?" Answer scale from 0% (extremely unlikely) to 100% (extremely likely). Confidence intervals are at the 90% level.

Figure 3.10: Elicited beliefs by income and score

divides the observations into above and below median scorers, casts doubt on this conclusion. As the right-hand side bar shows, the low average belief of low income subjects is strongly driven by above median scorers. By construction, the probability that an above median scorer with a low income receives the income corresponding to his or her relative performance must be higher in UP than in FI + R. This is because in 50% of the sessions included in UP, the earnings procedure is E, implying that a low income can be the result of bad luck in only about 50% of the cases. Yet, the mean belief of above median scorers in UP is lower than in FI + R, suggesting that one should be cautious when interpreting beliefs of this group.

Figure (3.10a) also hints at whether the results are due to underconfidence that is specific to the encryption task or whether a general self-effacing bias prevails under uncertainty. If participants are simply too underconfident regarding that particular task, the belief of high income subjects with a below median score should be lower than that of low income subjects with a below median score.[12] However, such a difference does not exist.

Recall that the second belief question, which is asked only in treatment UP, is about what participants think about the likelihood that earnings procedure E has been used in their session. Figure (3.10b) shows what participants infer about the earnings procedure of their session from their own outcome. As expected, the low belief of above median scorers with a low income translates into a low probability attached to E. All other sub-groups have intermediate beliefs about the probability of E, as implied by

[12] A low belief refers to a low probability attached to the event that relative income reflects relative performance.

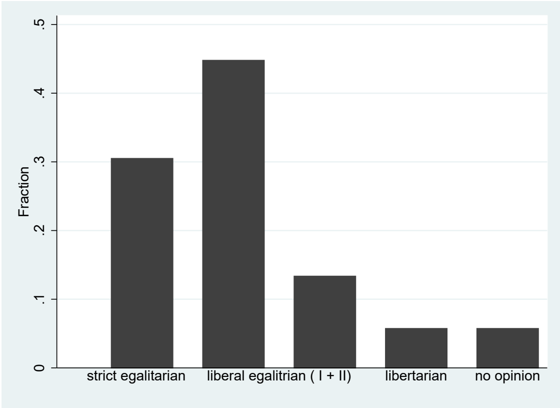

(a) Fairness ideals

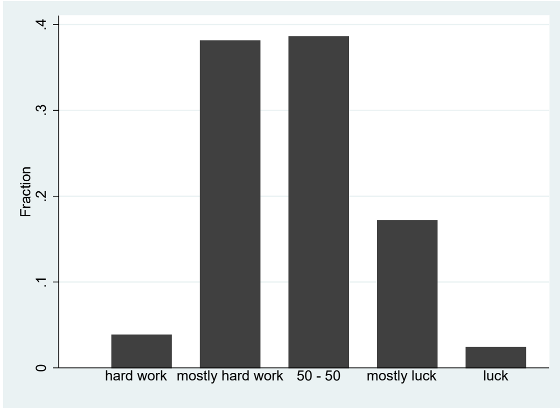

(b) Determinants of income

Figure (a) and (b) report responses from 210 participants in all treatments. Recall that the type I liberal egalitarian fairness ideal indicates a concern for equality of opportunity whereas the type II does not (see section 3.3.4). Confidence intervals are at the 90% level.

Figure 3.11: Distribution of survey responses

their beliefs about the influence of luck on their own outcome. Roughly speaking, these results support the claim that subjects were able to make the correct inference about the earnings procedure (R or E) in their session from their own outcome. Yet, the analysis of elicited beliefs is not as conclusive as the results obtained from the comparison of the main treatments. That is, it neither supports nor contracts the previous findings as the variance in the belief data seems to high.

3.4.6 Survey Responses

The distribution of survey responses is displayed in figure (3.11). Out of those who selected anything but “no opinion”, one third of respondents states a general fairness view that is in line with strict egalitarianism. About one half agrees with one of the two options in line with liberal egalitarianism whereas only about 10 percent selected the libertarian fairness ideal. This distribution shown in figure 3.11a roughly complies with the results obtained from the structural estimation by Cappelen et al. (2007) based on an experiment with Norwegian students.

According to figure (3.11b), about 80% of the participants believe that effort is at least as important as luck for earning a high income in Germany. This high value seems in line with existing research. In the 1992 ISSP, 87% of west Germans say that hard work is at least fairly important to get ahead in life. In summary, the survey responses suggest that the subject pool seems fairly standard with respect to both the prevalence of different fairness ideals and the perceived importance of luck for economic success in general.

3.5 Discussion

The heterogeneity in distributive preferences arising from different experiences of success is also present in Deffains et al. (2016) and Cassar and Klein (2017). By eliciting beliefs about the role of luck in the earnings procedure, the authors find that a self-serving bias explains the differences in distributive preferences in an effort-dependent environment similar to my treatment FI + E.[13] Their argument is that luck is inherent in almost any real-effort task, especially if random matching with an opponent player is involved. Hence, it is possible that participants form opposed beliefs about the role of luck affecting distributive preferences. Since I find strong evidence against the presence of such a self-serving bias, why then do distributional preferences persistently differ with income, even the earnings procedure is known to be random?

A promising explanation concerns the effect of experience and on the adoption of different fairness ideal as it may explain many of the results obtained. In treatment FI + R, differences in distributive preferences cannot be reasonably due to heterogeneous beliefs about whether inequality is due to factors that were under individual control, since people cannot be held accountable for (forced) bad luck. Instead, high income spectators might view the procedural fairness of treatment condition R differently than low income spectators. In particular, winning in this random earnings condition could make them attach more weight to the fact that opportunities were equal ex ante and thus make them prone to adopt a libertarian fairness ideal. In contrast, low income spectators pay less attention to ex ante equality of opportunities and are thus not willing to accept the resulting inequality. Importantly, heterogeneous perceptions of equality of opportunity are not due to a self-serving bias. That is, they do not result from participants attributing success to internal and failure to external factors. Instead, adjusting their perceived importance of equality of opportunity helps high income spectators to enjoy their own luck rather than feeling guilty for their opponent. Although the presented experiment does not test for the above mentioned mechanism explicitly, there is two pieces of evidence in the data supporting it.

First, the analysis of individual transfer decisions depicted in figure (3.5) in the previous section has shown that a large fraction of high income spectators in treatment FI + R distributed almost nothing, which is in line only with the libertarian fairness ideal. Second, the data from subjects' self-reported fairness ideals suggest a dependence on the experience of success. Figure (3.12) illustrated the suvrey data from both questions divided by income over different treatment combinations. I merge the data

[13]In the latter paper, the authors also note that a self-serving bias cannot reasonably explain all of the observed differences since the difference persists even when earnings are allocated randomly.

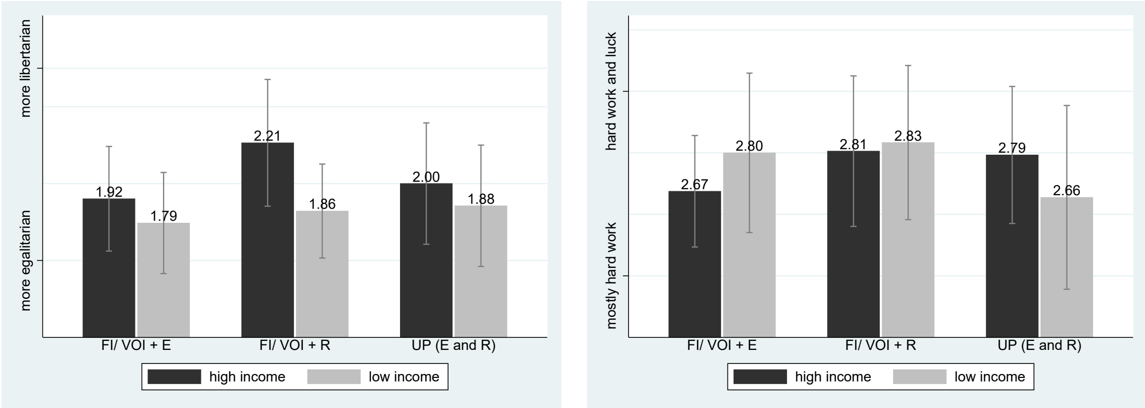

(a) Fairness ideals (b) Determinants of income

In figure (a), bar heights are based on the following nominal values assigned to survey responses: "strict egalitarian"=1, "liberal egalitarian (I)"= 2, "liberal egalitarian (II)"= 3, "libertarian"= 4. In figure (b), bar heights are based on "hard work"=1, "mostly hard work"=2, "hard work and luck"=3, "mostly luck"=4., etc.. Confidence intervals (at the 90% level) must be interpreted with caution as fairness ideal is (at best) an ordinal variable.

Figure 3.12: Survey responses by income

from VOI and FI in this figure because participants in the VOI treatments have already received information regarding their income when these questions are asked, rendering their received information participants have received in the FI and VOI treatments equivalent. Due to the focus on equivalent information for this part of the analysis, I distinguish between FI/ VOI + R and FI/ VOI + E but not between UP + E and UP + R as the latter two treatment combinations are characterized by the same information provided to participants as well. Notably, subjects' beliefs about the importance of hard work in real life are independent of income and treatment as shown by figure (3.12b). Since the underlying question explicitly addresses the situation in Germany, answers should indeed be independent of the income earned in the first phase of the experiment. Similarly, the prevalence of the libertarian libertarian fairness ideal should not vary with experience if individual fairness ideals were time-invariant. While this seems obvious when success is exogenous in treatment condition R, it should also be true in condition E. This is because the results depicted in figure (3.3) have shown that conditional on income, transfers are completely independent of the achieved score. However, figure (3.12) suggests that the fairness ideal of high income spectators tends more towards libertarianism than that of a low income spectators in all three treatment combinations. Surprisingly, the difference between high and low income subjects is even largest when randomness has determined income.[14] Hence, it seems that fairness

[14]Counting only individuals who state that they would always oppose redistribution (libertarian fairness ideal) yields a similar conclusion. In the random earnings condition, four high income subjects opt for the libertarian fairness ideal while it is selected only once by low income subjects.

ideals, rather than beliefs about the role of luck, are influenced through the experience of success and failure.

Most likely, the adoption of different in fairness ideals is not the only driver of the observed differences in distributional preferences. As already noted by Cassar and Klein (2017), in-group favoritism provides another potential explanation. In their influential paper, Chen and Li (2009) show that assigning subjects to groups merely on the basis of revealed preferences for Kandinsky pictures makes subjects more pro-social towards other in-group members. When people experience either economic success or failure as in the experiment presented, mutually exclusive income groups emerge as in the present experiment. As a consequence, people may act more pro-socially towards members of their own group, implying that subjects transfer less from the high to the low income subject if they have earned a high income themselves.

3.6 Conclusion

How does the experience of success affect distributional preferences under uncertainty differently than the experience of failure? Using an experiment in which the the earnings procedure is revealed to subjects only in some but not all treatments, I find a clear answer to this question: When participants do not know whether relative performance or randomness has determined income, distributional choices of successful (high income) spectators are identical to the choices of high income spectators when income is surely based on luck. Conversely, the choices of unsuccessful (low income) spectators are very similar to those chosen when earnings depend on performance. In general, the experience of success is associated with significantly lower levels of redistribution than the experience of failure. In addition, both successful and unsuccessful subjects are in favor of more distribution when income is assigned randomly than when it is based on relative performance. Hence, the diametrically opposed effect of experience on distributive choices under uncertainty implies that support for redistribution becomes more aligned when the source of inequality is unknown.

The second question this paper seeks to answer is why the experience of success and failure leads to different preferences over redistribution. Most importantly, the results described above reject the presence of a self-serving bias, which would make high income subjects believe that effort, rather than luck, has lead to their success. Instead, a self-effacing bias, which makes subjects attribute success to external factors such as luck and failure to internal factors such as poor performance, can explain the unambiguous conclusion from the participants' choices under uncertainty. Ruling out a self-serving bias as a source for the prevailing gap in distributional preferences between high and low income subjects calls for an alternative explanation. Interestingly, several findings from the presented experiment point towards the same direction. First, the distribution of individual transfers by high income spectators across all treatments shows that a large fraction of them redistributes very little or nothing, which is a choice that coincides exclusively with the libertarian fairness ideal. Second, a survey conducted at the end of the experiment reveals that self-reported, general fairness ideal are more libertarian among successful subjects. Hence, both findings suggest that fairness ideals are affected by experience, at least in the short-term.

Charitable giving depends to a large extent on the distributional preferences of those who have the means to give away parts of their wealth. Similarly, the power to distribute resources within organizations or the society as a whole often is in the hands of the most (economically) successful individuals. Consequently, the bias in distributional preferences induced by the defining experience of that group constitutes a major challenge to reaching a state in which resources are distributed just and fairly. A key finding from my experiment is that the bias in distributional preferences does not result from biased beliefs about the role of luck, but most likely from success-driven adoption of the libertarian fairness ideal. Hopefully, this insight will help to design instruments than can reduce the bias in the real-world.

4 Relative Earnings and Rank-Loss Aversion[1]

4.1 Introduction

Do people value relative income as a means or as an end? In parts, the answer to this question lies in the connection between income and status. More than a century ago, Veblen (1899) noticed that people may seek to signal high status by adopting certain behaviors. An important component of the resulting status expressing behavior constitutes the increased consumption of "positional goods" (Frank, 1985). Indeed, evidence suggests that this form of behavior leads to real material benefits. For example, Ball et al. (2001) find that status increases the chances of obtaining better economic outcomes even in a competitive market environment. Naturally, conspicuous consumption may signal high status only if it outweighs that of others. As a consequence, relative income is valued as a means but not necessarily as an end.

In contrast, others have emphasized the role of relative income without invoking the argument of relative consumption (Robson, 1992). If we subscribe to the view that people's preferences have been shaped by natural selection, we may conclude quickly that striving for better relative positions is an integral part of human nature and, thus, constitutes a means as well as an end. The reason is that in the major part of the course of human history, a better social standing has been associated with both higher survival rates and more descendants (Robson, 2001). At first glance, evidence in favor of such a general notion of rank concerns may seem widely available. Not only does happiness decrease in the income of those living close by (Blanchflower and Oswald, 2004), but also does job satisfaction in the income of one's peers (Card et al., 2012). While arguing that low relative income with respect to the entire society affects the ability to compete in conspicuous consumption seems plausible, it does not when

[1]I thank Henrik Orzen, Hans-Peter Grüner, Justin Leduc, Stefan Penczynski and Sander Renes for valuable comments.

referring to a few coworkers' income. Does this mean that income rank is valued in its own right? Not necessarily! In fact, evidence in support of a preference for higher relative income presented in many empirical studies may result from a strong aversion to disadvantageous inequality. For example, Card et al. (2012) are able to explain their results with an adopted version of the inequaity aversion model by Fehr and Schmidt (1999).[2]

Hence, existing evidence in support of pure rank-utility is highly inconclusive. Yet, gaining a full understanding of why and when people care about relative income is crucial for predicting the magnitude as well as the widespread of preferences for higher income ranks. In economics, accounting for such preferences is indispensable as they sometimes determine the limits to redistribution (Corneo and Gruner, 2000).[3] Therefore, this paper addresses the question whether relative income matters to people in the absence of any additional motives. More precisely, I report results from a laboratory experiment designed to test whether people refrain from giving up their (earned) income rank when maintaining the ranking neither reduces inequality nor has any status implications that could translate into material benefits. Instead, participants must trade off their income ranks with higher inequality while own income remains unaffected by their choice. That is, following an earnings phase, participants must distribute additional windfall money either to a poorer or richer participant, sometimes requiring them to give up their earned rank. Hence, the design allows to distinguish between two different explanations of rank concerns that usually point in the same direction.

The results suggest that rank-loss aversion affects some people's behavior, making them accept more inequality than in the absence of the rank-inequality trade-off. More precisely, I observe that when giving to the poorer recipient would require to give up one's rank, the average transfer to the poorer recipient is about 20% below the level prevailing in any other situation. While this decrease turns out to be significant only in some estimations, the effect on the number of choices leading to an overall increase of inequality always is. That is, I observe that when the earnings difference to the lower ranked participant is sufficiently low, 35% of participants facing such a situation give less than 50 % of the money to the poorer recipient, as compared to 7% to 21% in alternative situations. Here, estimation results confirm that the prevalence of this inequality increasing choice almost doubles from 17% to 30% when one's rank is at stake. In addition, the experimental design allows to investigate differences in rank-

[2]This includes the effect of no increase in job satisfaction when learning that one's peers earn less then oneself.

[3]The authors show that when signals of status are noisy, rank-utility effects may prevent desirable redistribution even if it is income rank-preserving.

loss aversion across several ranks as participants are matched into groups of six people with different endowments. Perhaps surprisingly, I do not find any evidence in support of abnormally stark rank concerns for particular positions in the income distribution.[4] This includes the second-to-last place even though subjects at this rank fall into the last place when giving up their rank, which people should be most averse to according to (Kuziemko et al., 2014).

At the beginning of the experiment, all subjects take part in a very brief estimation task in which they compete against other members of their group. The more precise a participant's guess relative to that of the other group members, the more she is allowed to earn in the subsequent real-effort task. Upon successful completion of the real-effort task, those with the best guess in the estimation task can earn up to 12 EUR. The second best guess allows to earn 10 EUR, the third best 8 EUR, the fourth best 6 EUR, the fifth best 4 EUR and the sixth best 2 EUR, respectively. Following the earnings procedure, participants are asked to repeatedly distribute additional money to two other members of their group. The trade-off between keeping one's rank and raising (not reducing) overall inequality arises because participants ranked second to fifth can distribute the additional money only to the subjects who are ranked one position below or one position above themselves. Asking individuals to distribute additional money consecutively for several times allows to investigate whether being in the position in which giving to the poorer recipient implies a rank loss leads to distributive behavior that significantly deviates from that in other situations.[5]

The experiment is designed such that the pattern of transfers chosen by rank-concerned participants differs from that of participants without rank-concerns would. I obtain those diametrically opposed predictions by integrating rank concerns into a modified version of the the model by Fehr and Schmidt (1999): Subjects who value their income rank simply as an end reduce the share they give to the poorer recipient when the difference in income between themselves and the lower ranked participant is small (but not too small). In some extreme cases, this implies that less than 50% of the money is given to the poorer recipient, thereby leading to an increase of the overall inequality. After the share distributed to the poorer recipient is reduced for one two consecutive rounds, it is likely to increase again once participants finally give up their rank in order to prevent a too stark increase in inequality. That is, rank loss aversion implies a non-monotonic relationship between the share that is distributed to

[4]The no-difference finding applies only to subjects ranked second to second-to-last as the experimental design does not permit to analyze rank concerns associated with the first and last place.

[5]Since total earnings of the recipients are accumulative, the actual number of repetitions is random and thus unknown to participants.

the poorer recipient and the earnings difference between him or her and the decision maker.

Different (regression) models consistently estimate the effect of rank-loss aversion to raise the frequency of giving less than half of the money to the lower ranked recipient from about 17 to at least 30%. Yet, the negative effect of rank-loss aversion on the absolute level of giving to the poorer recipient is estimated to be significant only for some specifications of the estimated model. Nevertheless, the consistently estimated significant increase in the number of people choosing to give less than 50% strongly suggests that rank-loss aversion influences some people's behavior, though it may be difficult to observe the effects on the aggregate level always.

4.2 Literature

This paper is related to an important strand of research that empirically investigates the relationship between income and happiness. Exploring survey data from the UK, Boyce et al. (2010) find that percentile in the income distribution predicts life satisfaction better than absolute income. Likewise, Blanchflower and Oswald (2004) and Luttmer (2005) find that own well-being is decreasing in the income of those living close. Using an experiment with hypothetical choice questions, Mujcic and Frijters (2012) observe that income rank matters irrespective of absolute income to Australian university students. Others have focused on smaller peer groups and investigated the relationship between happiness and income relative to one's peers. Clark and Oswald (1996) obtain results that are consistent with job satisfaction being dependent on relative but not on absolute income. More recently, Card et al. (2012) collected field experimental evidence showing that job satisfaction decreases when people learn that they are earning less than their peers. More generally, much of this literature contributes to the discussion of the Easterlin-Paradox (Easterlin, 1974), emphasizing that relative income, rather than absolute income, determines individual happiness.

As Heffetz and Frank (2011) note, the desire for relative income and status are closely related to the literature on social preferences. Seminal works in this area are from Fehr and Schmidt (1999), Bolton and Ockenfels (2000), Charness and Rabin (2002), and Engelmann and Strobel (2004). The link to this literature persists because in most cases, positional concerns can be explained by inequity aversion, and the aversion to disadvantageous inequality in particular. Indeed, Card et al. (2012) are able to show that their results match predictions obtained from a modified version of the inequity-aversion model by Fehr and Schmidt (1999). Importantly, this holds for both of their main findings: a decrease in job satisfaction after learning about peers with higher income and almost no effect on job satisfaction when learning about peers with lower income.

As already indicated, this paper contributes to the literature on status seeking behavior in more general. Beginning with Veblen (1899), many economists have argued that people value status and associated investments in conspicuous consumption goods because it allows them to achieve improved economic outcomes. Building on the closely related assumption that people have a desire for higher relative income, theoretical contributions such as Frank (1985), Robson (1992), Corneo and Gruner (2000) and Hopkins and Kornienko (2004) have explored the implications of status concerns for demand, inequality, growth and redistribution, respectively. The findings from this experiment lend credibility to these mostly theoretical works as they provide evidence in support of one of their key underlying assumptions.

Over the last decades, plenty of evidence in support of both status having positive effects on economic outcomes and people thus seeking for higher status has been documented. In Ball and Eckel (1996), Ball and Eckel (1998) and Ball et al. (2001), the authors manipulate status in the lab and observe that a higher status often leads to more favorable outcomes for those holding it. Rablen and Oswald (2008) find status effects outside the lab by estimating that winning a nobel price prolongs the winner's lifetime by one to two years. In the majority of the experimental research investigating the effects of status, the award of status is salient. As a consequence, the awardee of high status is treated differently by others in subsequent social interactions, which is often beneficial for the awardee. Instead, this papers ask the slightly different question whether people seek higher income ranks in the absence of any public announcements. Even though this is the case in many real world applications, evidence showing that people care about private income ranks as they care about publicly awarded status is still scarce.

The work by Dijk et al. (2014) is among the very few papers trying to answer whether pure income rank concerns affect behavior. Comparing portfolio investment decisions in the lab, the authors observe that information about peers' performance induces investment choices that compare with those made under a tournament incentive scheme. Hence, they conclude that rank matters even in the absence of additional material incentives. In Kirchler et al. (2018), the authors bring the lab into the field, confirming most of the evidence that peer comparisons induce more risk-taking, especially among professional investors. Another work belonging to this area is the work by Kuziemko et al. (2014), which also closest to my paper in terms of the experimental design. In a series of experiments and surveys, the authors collect evidence in support of last-place aversion, which can be seen as a rank-specific status or rank concern. Yet, their identification relies on inter-rank comparisons, preventing them from learning anything about an aversion to give up one's rank across all ranks. Moreover, rank in their experiment is assigned randomly wheres it is associated with some merit in this.

More broadly, this paper is also related to the experimental literature on distributional preferences in large groups. While the research question investigated by Durante et al. (2014) or Erkal et al. (2011) is different from mine, rank-loss aversion is likely to influence individual behavior in their experiments as well, thereby changing how their data can and possibly even should be interpreted. For example, selfishness is not the only explanation for the relatively low transfers chosen by participants ranked first in Erkal et al. (2011). Instead, participants ranked first may fear a loss of their rank most when choosing a transfer that is too high even without being most rank-loss averse. This is because they can be almost certain that they will not receive any transfer from other members of their group, while any other member may receive up to several transfers.

To summarize, the existing literature has focused on when and why people value status and relative income. As the present papers addresses the lack of evidence in support of pure income rank concerns, its major contribution lies in this area. In addition, though to a much lesser extent, I contribute to the literature on the behavioral implications of status. A common finding from this mostly experimental research is that those with an advantageous position are the least helpful and most greedy (Piff et al. (2010), Guinote et al. (2015), Piff et al. (2012)). Using a dictator game to measure social preferences of different status groups, Fisman et al. (2015) also find that high status subjects are least pro-social. The experimental findings presented in this paper complement this line of research by showing that people at any rank, not only those with a high status, may act less pro-social if it enables them to preserve the existing ranking.

4.3 Theory and Experiment

4.3.1 Experimental Design

The experimental design aims at identifying the effects of rank-loss aversion. To generate a ranking, the experiment uses a real-effort task (first stage). To detect potential effects of rank-loss aversion, the experiment forces participants to trade off their rank with reducing overall inequality in a distribution task (second stage).

At the beginning of the first stage, subjects are randomly matched into groups of six and then perform a brief estimation task to determine a preliminary ranking.[6] As a part of the estimation task, subjects are shown a large table containing 1,200 numbers which are either one or zero for 15 seconds. Afterwards, they are asked to provide an estimate of the number of zeros in the table that they have seen. Based on the relative precision of the estimate within their group, subjects are assigned tasks of varying difficulty for the subsequent real-effort phase.[7] The assignments are fixed across all groups and depend only on relative performance. If all subjects successfully complete their task, those facing the more difficult task also achieve higher earnings and the preliminary ranking equals the final ranking.

The real-effort task is based on Abeler et al. (2011). Subjects are required to count the number of zeros in as many tables as their assignment prescribes (see figure (C.2) in appendix C). If subjects count correctly, they are rewarded with 0.50 EUR per table. Subjects ranked first after the estimation task may count up to 24 tables. From the second to sixth rank, the assigned upper limits are 20, 16, 12, 8, and 4 tables. All subjects, irrespective of their rank, are given 20 minutes to complete this task. Fulfilling the assignment is voluntary and subjects are free to count fewer tables, but not more. Hence, a subject ranked first in the estimation task can at most earn 12 EUR, a subject ranked second at most 10 EUR and so forth. The total time of 20 minutes is chosen such that almost all participants are able to reach their implied earnings cap, though it certainly requires higher effort from those at higher ranks. When a subject completes her assignment before the time is over, her screen switches to an article about Mannheim, which is copied from the English Wikipedia. Prompting participants to read the article is simply a measure to keep them focused and concentrated.

During the second stage, subjects decide on how to divide additional windfall gains between two others members of the same group. To create the afore mentioned trade-off, the two possible recipients are the subjects ranked one position above and one position below the subject making a decision if she is ranked second to second-to-last. That is, someone ranked second can distribute to the participants ranked first or third, someone ranked third to the participants ranked second or fourth, and so forth. Importantly, a subject's rank in this second stage refers to the income ranking obtained after the real-effort phase. By construction of the experiment, the choice sets of subjects ranked first and last have to be different. If ranked first, subjects must distribute the additional money between the group member ranked second and the member ranked last. If ranked last, subjects distribute between those ranked first

[6]Each subject in a group is assigned a (group) unique ID from the list Blue, Brown, Green, Orange, Red, and Yellow. In the on-screen instructions, participants are henceforth referred to as player [color].

[7]If a tie occurs, the computer allocates the higher rank randomly.

and fifth in their own group. Despite the resulting disadvantage of not being able to explore the degree of rank-loss aversion among subjects ranked first and last, assigning systematically different recipients to those subjects also has the advantage of creating a control group as they never face the inequality-rank trade-off.

Distributing additional money takes place over several, consecutive rounds, with transfers accumulating over rounds. While the instructions only mention that the number of rounds is determined randomly, it is, in fact, five plus a random number drawn from a Poisson distribution with mean equal to one. In each round, the additional money that must be distributed between the two recipients amounts to 1.00 EUR. Using a slider, participants can select the share they wish to distribute to each of the other participants with a precision of 0.01 EUR (see figure (C.3) in appendix C). Besides, the screen shows a large table containing information about all group members income during the entire course of this task. Importantly, the figure showing the post-distribution income (see column "total income" in figure (C.3) in appendix C) is updated immediately whenever the slider is moved by the participant. Likewise, the ranking positions are adjusted instantly in response to a change of a subject's income rank.

Because subjects never know whether there will be another round, they should never withhold money from a recipient with the intention to compensate him or her in the subsequent round. Instead, they should always distribute the money such that the resulting accumulated earnings allocation matches their (constrained) most preferred distribution. After some rounds, the income difference with respect to the poorer recipient becomes small, resulting in the intended trade-off requiring subjects to choose between reducing inequality further or preserving the ranking. At the end of the distribution task, all decisions of a single participant per group are implemented. This rules out strategic behavior and permits to interpret the observed decisions as the expression of individual preferences.

Upon completion of the distribution task, all subjects are asked to participate in a simplified trust game prior to learning the outcome of the distribution phase. Participants are informed about the existence of this last part at the beginning of the experiment. Yet, they receive no information about this part except that the major share of their earnings from this experiment will be determined during the first part.[8]

A comment on the design.

The earnings scheme resulting from the tournament to allocate earnings implies that differences in income between adjacent ranks are constant. As a result, distributional

[8]In the trust game, I measure heterogeneity in trust and trustworthiness across income ranks.

choices of subjects are not only comparable to subjects at identical ranks in other groups, but to other subjects within the same group as well. While constant payoffs are a major advantage of using a tournament, a standard tournament typically leads to several issues that would be detrimental to identifying the consequences of rank-loss aversion. Below, I discuss those issues in more detail and explain how the chosen design resolves them.

First, if the differences in income determined by a fixed-prize tournament are not proportional to the differences in performance, participants are likely to be malevolent towards higher ranked participants (Grund and Sliwka, 2005). The resulting enviousness or guilt from the perspective of those winning the higher prize may significantly affect distributional preferences. As differences in performance (and the underlying effort) are naturally small for those with an average position, this described type of malevolence or guilt would therefore systematically distort the experiment's key trade-off between the preference for an equal distribution and rank-loss aversion.

Second, when evaluating inequality, subjects may not only pay attention to others' income, but also to the effort provided to earn it. As Abeler et al. (2011) point out, the counting task certainly entails a positive cost of effort since it is very tedious. According to the theory described in section 4.3.2, costs of effort influence a participant's distribution decision if she takes effort into account while evaluating the degree of inequality. Assigning tasks of varying and predefined length to participants based on a very brief estimation task makes the effort provision roughly proportional to earned income. It thus constitutes another measure to improve comparability across ranks.

Third, using the estimation task to indirectly determine final income while the vast majority of effort must be provided in another task separates skill and talent in the real-effort task from the attained income rank. That is, subjects ranked second will not necessarily find the real-effort task easier than subjects ranked second-to-last since the second place does not result from being second best in the real-effort task. Separating skill from income is again crucial when relying on the comparability of the rank-inequality trade-off. Due to the potential connection between income and effort, the rank-inequality trade-off would be different across ranks if the perceived cost of effort was. Separating income from skill that facilitates earning the income prevents this distortionary effect as well as other endogeneity problems typically resulting from self-selection.

4.3.2 Theoretical Framework

Suppose that an ordering $y_1 < y_2 < ... < y_6$ exists regarding the income levels of all members of a group. Further, assume that utility is additively separable in "standard utility" and rank utility. Denote by R_i individual i's actual rank after redistribution and by R_i^e her income from the real-effort phase. A simple way to incorporate rank-loss aversion into individual i's preferences is to write her utility as follows:

$$\Gamma_i(\gamma, y_i) = U_i(\cdot) + G(R_i^e, R_i) \quad (4.1)$$

where $U_i(\cdot)$ captures i's "standard" social preferences neglecting rank-loss aversion and where the additional term $G(R_i^e, R_i)$ captures utility from her final rank R_i, which depends on her earned rank R_i^e as well. The general idea of combining the preference for a particular rank with existing models of social preferences as in (4.1) is borrowed from Kuziemko et al. (2014). Nevertheless, I deviate from Kuziemko et al. (2014) in three ways. First, I assume that not only people in the second-to-last but in any place are rank-loss averse. Second, I account for the recent evidence on fairness views by assuming that people care not only about self-centered inequality, but also about inequality between others, which helps to derive more reasonable predictions.[9] Finally, the differences in design between the distribution phase in Kuziemko et al. (2014) and this experiment require to derive new predictions that account for several consecutive distribution rounds.

In the experiment, participants first attain a particular rank according to their performance and can then choose actions that either preserve their rank or make them fall behind by one position. Hence, the experiment explores the change in utility from giving up a position, but does not measure absolute utility from rank. Thus, testable predictions rely only on the difference in those utilities, which I denote by R':

$$R' = G(R_i^e, R_i^e) - G(R_i^e, R_i^e - 1) \quad (4.2)$$

Even though this assumption can be relaxed, for simplicity assume that R' is constant across all ranks. Naturally, the functional form assumptions regarding $U_i(\cdot)$ influence the predictions from any model of rank-loss aversion, even more so when accounting for preferences over inequality between others. To provide a general overview of predictions that could obtain, I first discuss those resulting from three of the most prominent social preference models.

[9] A motivation for this assumption is given later in this section.

According to the equity-reciprocity model by Bolton and Ockenfels (2000), utility depends on two factors: own income y_i and the individual share of the total surplus $\bar{y}$. Since the design keeps own income y_i constant while not allowing participants to influence $\bar{y}$ (the decision maker must distribute all of the additional money), the model predicts that participants would always be indifferent between giving to the richer or poorer recipient. With rank-loss aversion instead, participants would always give everything to the richer recipient if this is necessary to secure R', the additional utility from keeping her earned rank.

Similarly, Charness and Rabin (2002) suggest that taking into account preferences over $\frac{y_i}{\bar{y}}$, y_i and the income of the poorest player is sufficient to explain individual behavior in distribution games. By design, the second-to-last ranked player is the only participant who can give to the poorest (last-ranked) subject. This means that all but the participant ranked second-to-last, should be indifferent between giving to the richer or poorer recipient.[10] Yet, when accounting for rank-loss aversion, predictions are identical for all subjects, including those ranked second-to-last. The reason why someone ranked second-to-last will give everything to the richer recipient if necessary to secure R' is because she becomes the poorest member of the group herself by giving up her rank. When this happens, she cannot make the poorest member any richer and the incentive to give up her rank vanishes.

The third model is the one by Fehr and Schmidt (1999), which again leads to quite similar predictions. The authors assume that utility of individual i is given by:

$$U_i(y_i) = y_i - \alpha_i \frac{1}{n-1} \sum_{j \neq i} \max\{y_j - y_i, 0\} - \beta_i \frac{1}{n-1} \sum_{j \neq i} \max\{y_i - y_j, 0\}$$

where $\beta_i \leq \alpha_i$, $0 \leq \beta_i < 1$ and y_i represents i's monetary payoff. The model predicts that a decision maker should strictly prefer giving to the poorer of the two recipients if the poorer recipient's income puts him in the domain of advantageous inequality. Otherwise, she should be indifferent between giving to either subject. Accounting for rank-loss aversion makes the indifference disappear. Instead, the prediction is that participants would give up to 100% of the additional money to the richer recipient to prevent the loss of R'.

Notably, the models discussed were developed to mainly explain behavior in games which require people to trade off own material interests with giving a fair share to others. As a consequence, it was sufficient to build these models on purely self-centered

[10] In fact, the last-ranked subject is included in the first ranked subjects' choice sets as well. Yet, a subject ranked first never faces a trade-off between giving to the poorer recipient or maintaining her rank in this experiment. Thus, the prediction is to give to the poorer recipient always.

other-regarding preferences, that take into account only comparisons between others and oneself.[11] Instead, the experiment requires participants to be spectators and choose a distribution of money between other recipients without being able to give anything to themselves. Besides, the prediction of distributing all of the additional money to the richer recipient, even when R' is arbitrarily small, seems quite implausible and not very robust in the light of the existing experimental evidence on general fairness views. Specifically, Cappelen et al. (2007) as well as Cappelen et al. (2010) have found that people also care about how wealth is distributed among others. That is, they tend to accept inequality between others if they hold them accountable for the outcome but support redistribution to reduce the inequality otherwise. To account for such general fairness concerns, I add an additional term capturing disutility from inequality between others to the framework by Fehr and Schmidt (1999). This extended version of the inequity-aversion model writes:

$$U_i(y_i) = y_i - \alpha_i \frac{1}{n-1} \sum_{j \neq i} \max\left\{y_j - y_i, 0\right\} - \beta_i \frac{1}{n-1} \sum_{j \neq i} \max\left\{y_i - y_j, 0\right\} \tag{4.3}$$

$$- \frac{1}{n-1} \gamma \sum_{j \neq i}^{n} \left(\sum_{k \neq i,j}^{n} \max\left\{y_k - y_j, 0\right\} \right) \tag{4.4}$$

where $0 < \gamma < 1$ is the weight attached to the cost arising from inequality between others. The linearity and additive separability of the utility function as well as the experimental design simplify predicting behavior in the distribution task significantly. If there was a fourth participant who was in the domain of disadvantageous inequality with respect to the richer recipient but in the domain of advantageous inequality with respect to the poorer recipient, the disutility from that additional inequality would matter for the decision of how much to give to the poorer and richer recipient as well. Yet, all other subjects are either richer than both or poorer than both recipients as the recipients the money can be distributed to are ranked exactly one position above or below the spectator. Consequently, the shape of the utility function implies that how the additional money is distributed between the two recipients does not affect utlity from inequality with respect to any other participant. Therefore, it is sufficient to consider only the spectator's income and that of the two recipients when deriving predictions for the experiment's distribution task.

Consider an individual i with income y_i and an earned rank $R_i \in \{2, 3, 4, 5\}$. Denote the poorer recipient's earnings at the beginning of any round by $\underline{y}$ and the richer recipient's earnings by $\bar{y}$, respectively. Let $\delta \in [0, 1]$ be the predicted amount given to the poorer recipient. Further, denote by $\Delta \in \mathbb{R}$ the difference in total income between

[11] The model by Charness and Rabin (2002) is an exception as the wealth of the poorest recipient matters independetly of own income. Yet, the design still leads to identical predictions in this case.

subject i, who chooses a transfer δ, and the poorer recipient at the beginning of a round. Then, i keeps her rank R_i and does not lose R' iff

$$\underline{y} + \delta \leq y_i \Leftrightarrow \delta \leq y_i - \underline{y}. \tag{4.5}$$

The constraint (4.5) is binding if the amount i would choose to give to the poorer recipient in the absence of rank concerns, denoted by δ^*, violates (4.5). To investigate the trade-off that arises when (4.5) is binding, I derive the optimal choice of δ^* first. From (4.3), the solution to δ^* can be obtained as follows:[12]

$$\delta^* = \arg\max_{\delta\in[0,1]} \Big(-\alpha\big[\max\{\bar{y} + (1-\delta) - y_i, 0\} + \max\{\underline{y} + \delta - y_i, 0\}\big] \\ -\beta\big[\max\{y_i - (\underline{y} + \delta), 0\}\big] - \gamma\big[\max\{\bar{y} + (1-\delta) - (\underline{y} + \delta)\}\big]\Big)$$

For $\Delta \geq 1$, i always gains $(\delta \cdot \beta)$ from giving the entire share $(\delta = 1)$ to the poorer recipient as it reduces advantageous inequality and minimizes the additional loss from an increase in disadvantageous inequality $(1-\delta)\cdot\alpha$. For $\Delta < 1$ instead, i's disutility from disadvantageous inequality does not depend on how the amount of $(1-\Delta)$ is distributed between the poorer and richer recipient. Yet, the choice of δ still affects i's (dis-)utility from inequality between others according to $-\gamma(1-2\delta)\ \forall\ \delta < \frac{\bar{y}-\underline{y}+1}{2}$. Therefore, i is not indifferent between giving to either recipient for $\Delta \leq 1$ (and even $\Delta \leq 0$) but prefers to reduce the persisting inequality further. Once equality between the two recipients has been reached, i splits the money equally to prevent additional inequality.[13] Formally, subject i chooses:

$$\delta^* = \min\{\bar{y} - \underline{y}, 1\} + \frac{1}{2}\big(1 - \min\{\bar{y} - \underline{y}, 1\}\big) \tag{4.6}$$

$$= \frac{1}{2}\big(1 + \min\{\bar{y} - \underline{y}, 1\}\big)\ \forall\ \bar{y} \geq \underline{y} \tag{4.7}$$

Notice that the transfer given in (4.7) accounts for preferences over all domains of inequality since $\Delta \geq 0 \Rightarrow\ \bar{y} - \underline{y} > 1$, implying that whenever there is advantageous inequality gives rise to a higher transfer, inequality between others does as well. However, δ^* is the optimal choice only in the absence of rank-loss aversion. If instead, i's preferences exhibit rank loss-aversion as suggested in (4.1), a trade-off arises that sometimes implies a different choice. Importantly, a different choice occurs only if (4.5)

[12] The constant y_i as well as the standardization by $1/(N-1)$ are omitted as they do not affect i's choice.

[13] Since people have exerted effort to earn their income, people may consider some part of the persisting inequality as justified and may not seek to reduce it until full equality prevails. This possibility is explored in the next subsection.

is binding as an individual need not give up R' and can simply benefit from reducing both advantageous inequality and inequality between others otherwise. Conditional on preserving the ranking, i still wants to reduce inequality as much as possible. Thus, the optimal rank-preserving choice $\tilde{\delta}$ satisfies $\tilde{\delta} = \min\{\max\{y_i - \underline{y}, 0\}, 1\}$. Hence, $\delta^* = \tilde{\delta} = 1$ if $\Delta = y_i - \underline{y} \geq 1$ but $\delta^* > \tilde{\delta}$ otherwise. Then, for $\Delta < 1$, i chooses $\delta = \tilde{\delta}$ instead of δ^* if not giving up her rank outweighs the differences in utility from inequality between others. Formally, the condition writes:

$$\gamma\Big(\underbrace{\big[\bar{y} + (1-(y_i - \underline{y})) - \big(\underline{y} + (y_i - \underline{y})\big)\big]}_{\textit{inequality given } \tilde{\delta}} - \underbrace{\big[\bar{y} + (1-\delta^*) - (\underline{y} + \delta^*)\big]}_{\textit{inequality given } \delta^*}\Big) \leq R'$$
$$\Leftrightarrow 2\big(\delta^* - (y_i - \underline{y})\big) \leq R'/\gamma$$

By (4.7), $\delta^* = 1$ if $\bar{y} - \underline{y} \geq 1$, and $\delta^* = 1/2 + 1/2(\bar{y} - \underline{y})$ otherwise. Substituting δ^* into the above inequality yields Proposition 4.1.

Proposition 4.1 *The amount given to the poorer recipient δ depends on Δ as follows:*

$$\delta = \begin{cases} 1 & \textit{if } \Delta \geq 1 \\ \underbrace{y_i - \underline{y}}_{=\tilde{\delta}} & \textit{if } 1 > \Delta \geq 1 - \dfrac{R'}{2\gamma} \\ \underbrace{1/2 + 1/2\big(\min\{\bar{y} - \underline{y}, 1\}\big)}_{=\delta^*} & \textit{if } \Delta < 1 - \dfrac{R'}{2\gamma}. \end{cases} \tag{4.8}$$

By proposition 4.1, δ depends on $\bar{y} - \underline{y}$ for $\Delta < 1 - \frac{R'}{2\gamma}$. Substituting $\underline{y} = y_i - \Delta$ into the third line of (4.8), I obtain that

$$\delta^* = \begin{cases} 1 & \text{if } \Delta \geq 1 + y_i - \bar{y} \\ 1/2 + 1/2(\bar{y} - \underline{y}) & \text{if } \Delta < 1 + y_i - \bar{y} \end{cases} \tag{4.9}$$

The step-wise characterization of $\delta(\Delta)$ shows that the amount given to the poorer recipient is largest for $\Delta \geq 1$ and $\Delta \in [1 + y_i - \bar{y}, 1 - \frac{R'}{2\gamma})$. Moreover, i may never want to give up her rank but choose $\delta = \tilde{\delta} = \Delta$ for any $\Delta \in [0, 1]$ if the degree of rank-loss aversion is too strong in the sense that $R' \geq 2\gamma$. This is because in that case, the condition $\Delta \geq 1 - \frac{R'}{2\gamma}$ stated in (4.8) holds always. Yet, if i reaches $\Delta < 0$ at some point, it means that her degree of rank-loss aversion satisfies $R' < 2\gamma$. That is, conditional on observing choices for $\Delta < 0$, I obtain that:

Corollary 4.1 *$\delta(\Delta)$ is non-monotonic in Δ and has a local minimum at $\Delta = 1 - \frac{R'}{2\gamma}$.*

4.3.2.1 Accounting for effort

Eventually, participants assess inequality and relative standing in terms of both people's income and the effort exerted to earn it. That is, subject i might care about how her income (y_i(net of the cost of effort (c_i) compares with any other subject's net income $y_j - c_j,\ j \neq i$. Formally, this can be accounted for by substituting $y_i - c_i$ for y_i into the equations (4.3) to (4.9) derived above. Since the difficulty of the task is chosen such that all subjects deliberately complete their assignments, I assume that $y > c$ for all ranks. Besides, the proportional relationship between the the number of correctly counted tables and earned income suggests that cost of effort is a proportional function of y_j. [14] Thus, assume that $c_j(y_j) = a \cdot y_j,\ a \in (0,1)$, implying that that the degree of inequality between i and j given by $(1-a)(y_i - y_j)$. Overall, all threshold levels stated in proposition (4.1) have to be multiplied by $1/(1-a)$, which raises the critical level of Δ where giving is predicted to be minimal. To see why, note that giving is lowest if $(1-a)(y_i - \underline{y}) = 1 - R'/2\gamma$. By solving for $\Delta = y_i - \underline{y}$, I immediately obtain the following result:

Corollary 4.2 *If costs of effort matter to an individual's assessment of inequality, δ is minimal at $\Delta = \frac{1}{1-a}(1 - R'/2\gamma) > (1 - R'/2\gamma)$.*

The prediction of corollary 4.2 is rather intuitive. If costs of effort reduce the perceived inequality from differences in income, the poorer recipient, who has exerted less effort but whose earnings after having received additional money are almost as high as those of subject i, may even be considered better off than i. Thus, one might want to argue that in such a case, the higher rank is already lost to the poorer recipient despite a still positive income gap. Consequently, δ is chosen to be lowest when the decision maker is still better off including her higher incurred costs of effort, which coincides with a larger income gap than in the absence of any costs of effort.

4.3.3 Identification and Hypotheses

Some additional assumption regarding the magnitude and distribution of the rank-loss aversion parameter R' as well as the inequality aversion parameter γ are necessary in order to obtain testable implications from the model derived in section 4.3.2. Since the degree of rank loss-aversion may vary substantially across individuals, the first objective is to find the interval that captures most people's critical difference $\Delta = 1 - R'/(2\gamma)$ where the amount given to the poorer recipient is predicted to be lowest. The main

[14] The analysis of the data presented in section 4.4 will confirm the simplifying assumptions made here. That is, almost all participants count as many tables as possible and total counting time, which seems to be a more accurate measure of effort, is almost proportional to earned income as well.

part of the empirical analysis relies on the assumption that this critical interval is given by $[0.25, 0.75)$. The subsequent paragraphs provide the underlying arguments for this choice. Moreover, I present results from an alternative specification of the critical interval in the robustness check in section 4.5

Note that the value of γ, which measures disutitlity from inequality between others, should not exceed the values of α or β for most participants. Otherwise, an individual would be more concerned about inequality between others than between others and him or herself, which seems unlikely. Besides, most experimental results confirm that people dislike disadvantageous inequality more than advantageous inequality, implying $\beta < \alpha$. Fehr and Schmidt (1999) show that a β which is distributed between 0 and 0.6 with mean 0.31 can explain individual behavior across multiple games. Thus, a value of $\gamma = 0.25$, which implies that disutility from inequality between others is slightly lower than disutility from advantageous inequality appears reasonable. Concerning the distribution of R', I assume that $2\gamma > R'$ constitutes an upper bound since an individual would never want to give up her rank in the experiment otherwise. As there may as well be some people who do not care about their rank, suppose that $R' \geq 0$ constitutes the lower bound of the distribution of R'. Consequently, the critical threshold, $1 - R'/2\gamma$, where the transfer to the poorer recipient is minimal should be distributed over the interval $[0, 0.5)$. Importantly, this interval is not yet the interval in which I expect to *observe* the lowest transfers by most individuals for two reasons.

First, the threshold $\Delta = 1 - R'/2\gamma$ is to be multiplied by $1/(1-a)$ where a is the relative cost of effort to obtain one unit of y whenever costs of effort enter inequality concerns. In Abeler et al. (2011), subjects find it worthwhile to exert the effort necessary to count the number of zeros for 10 cents per table. In my experiment, the effort required to count all zeros is higher because I chose the ratio between zeros and ones is to be more balanced.[15] Therefore, I assume that $a = 1/3$, expanding the critical interval to $[0, 0.75)$. Second, the theory predicts not only that giving is lowest at $\Delta = 1 - R'/2\gamma$, but also that it is increasing over the entire interval $[1 - R'/2\gamma, 1]$. Due to the endogeneity of Δ, very few participants will in fact make a decision when being exactly at their critical threshold. Rather, most observations that are affected by rank loss-aversion belong to the entire interval $[1 - R'/2\gamma, 1]$.

Hence it is advisable to include observations with larger values of $\Delta > 1/(1-a)(1 - R'/2\gamma)$ as well by extending the interval to $[0, 1)$. Yet, an interval that is too large will almost surely contain both low and high transfers to the poorer recipient from the same individual. To see why, notice that for a particular level of individual rank-loss

[15] The size of the number tables is identical in Abeler et al. (2011) and my experiment (150 randomly ordered zeros and ones). Yet, the mean probability that a zero is drawn amounts to roughly 50% in my experiment, while it equals 30% in theirs.

aversion, almost every participant will at least once not be affected by the inequality-rank trade-off despite $\Delta \in [0, 1)$. This is because the income difference Δ is either still too large for Δ smaller but close to one, or too small, implying that the rank is already lost for Δ larger but close to zero. As a consequence, the chance of detecting an effect of rank-loss aversion would be mitigated through the use of an interval that is too large. Relying on the smaller interval $[0.25, 0, 75)$ is an approach to reduce this problem.

To allow for a reasonable comparison between giving behavior within the critical interval and that within other regions of Δ, I construct additional intervals of identical width outside of $[0.25, 0.75)$. In the regression analysis, dummy variables indicate whether the choice of δ belongs to a particular interval D_j where $D_1 = 1$ if $\Delta \in [1.25, \infty)$, $D_2 = 1$ if $\Delta \in [0.75, 1.25)$, $D_3 = 1$ if $\Delta \in [0.25, 0.75)$, $D_4 = 1$ if $\Delta \in [-0.25, 0.25)$, $D_5 = 1$ if $\Delta \in [-0.75, -0.25)$, $D_6 = 1$ if $\Delta \in [-1.25, -0.75)$, $D_7 = 1$ if $\Delta \in [-1.75, -1.25)$, $D_8 = 1$ if $\Delta \in (-\infty, -1.75)$ and $D_j = 0 \; \forall \; j$ otherwise.

To explore the effect of rank loss-aversion ($D_3 = 1$), I focus on two outcomes that can be computed from the available data. The first variable is called *min_giving*: Instead of taking into account all decisions, *min_giving* takes on the value of the lowest amount given to the poorer recipient within each interval. This method successfully deals with the problem that arises when multiple decisions are observed within $[0.25, 0.75)$. That is, based on the theory derived in section 4.3.2, I expect participants to give less than usual to the poorer recipient first and then decide to give up their rank in the next round, distributing a significantly higher share to the poorer recipient again. Unless an individual's threshold of $1 -' /2\gamma$ is smaller than 0.25, it is possible that both decisions are made when $\Delta \in [0.25, 0.75)$. Focusing on the lowest transfer is a way to handle those misleading observations without changing the interpretation of any potential effect for $D_3 = 1$.

In addition, I consider the relative frequency of giving less than 50% of the money to the poorer recipient such that inequality is, as a consequence, increasing. The corresponding outcome *less_than_half* is a dummy variable that takes a value of one if a subject chooses $\delta < 0.5$ in any given interval at least once and a value of zero otherwise. Counting only one event per individual and interval is again crucial as potential consequences of rank-loss aversion can hardly be detected otherwise for the same reason justifying the focus on the lowest transfer per interval when considering *min_giving*.

Notably, it is difficult to find a reasonable explanation for why participants whould give less than 50% to the poorer recipient without accounting for rank-loss aversion. Hence, an increase in the frequency of such choices would provide evidence that some people are rank-loss averse. Additionally, focusing on the extreme outcome of *less_than_half*=1 may be a more effective way of detecting the consequences of rank loss aversion. Simply because the task asks subjects to select any amount "between" 0 EUR and 1 EUR, most of them are going to distribute at least some positive amount to the richer recipient as well, implying $\delta < 1$ across all intervals. If there is only some additional noise in the decision making process, identifying a statistical difference between the averages of *min_giving* when $\Delta \in [0.25, 0.75)$ and *min_giving* when $\Delta \notin [0.25, 0.75)$ becomes incredibly hard unless the number of observations is very large. In contrast, the theory unambiguously predicts that people should be giving less than 50% only if it is necessary to maintain their rank, but never otherwise. Moreover, it seems unlikely that many participants make a "mistake" and sometimes choose a share smaller than the salient threshold of 50% by accident. Hence, a small number of observations with subjects giving less than 50% for $\Delta \in [0.25, 0.75)$ may be sufficient to identify an effect of rank loss-aversion when it should never be observed for $\Delta \notin [0.25, 0.75)$.

Notice that the experiment generates multiple observations from each individual such that each round of making a distribution decision corresponds to a new period. To account for the implied panel structure of the experimental data, variables are indexed by t in the following regression models. That is, $D_{j,}^{it}$ corresponds to subject i's decision in round t.[16] An important advantage of the panel structure is that it allows to include individual fixed-effects. Since individual fixed effects control for all time-invariant characteristics of an individual such as gender, social background or education, those variables need not to be included in the regression analysis. The first estimation builds on the following individual fixed-effects model:

$$Y_{it} = \alpha_i + \sum_{j \neq 3} \beta_j D_j^{it} + \epsilon_{it} \tag{M1}$$

where ϵ_{it} is an idiosyncratic error term. The dependent variable Y_{it} is replaced either by *min_giving* or by the dummy variable *less_than_half*. Accordingly, the estimated coefficient for β_j can be interpreted either as the predicted absolute change in the lowest transfer (to the poorer recipient) per decision maker or as the change in the relative

[16]The outcome variable Y_{it} is not only a panel variable, but also highly auto correlated. This is because $min\_giving_{i,t-1}$ in period $t-1$ determines Δ_t in period t and, thus, D_j^{it} $\forall$ j. However, the design rules out any causal effect of $Y_{i,t-1}$ on $Y_{i,t}$ as uncertainty about the total number of rounds incentivizes subjects to decide in any round as if the current round is the last round. That is, when choosing how to divide the additional money in the first round, subjects select the allocation they would be most satisfied with if the experiment ended afterwards. In the following round(s), subjects do not have to compensate a recipient for a low transfer in any previous round as that transfer was simply just. Consequently, what matters to subjects when distributing the additional money in round t is not the total amount already distributed but what the final allocation resulting from the decision in round t would be. The perfect collinearity of Δ_{it} and the final allocation in round t for any transfer δ_t ensures that the effect of the latter is captured by the use of the interval dummies D_j.

frequency of giving less than 50% when $\Delta \in [0.25, 0.75)$. For $Y_{it} = min\_giving_{it}$, the formal hypothesis writes:

$$\mathbf{H_0} : \beta_j \leq 0 \ \mathbf{vs.}\ \mathbf{H_1} : \beta_j > 0 \ \forall \ j \neq 3$$

while it is reversed for the outcome *less_than_half*. This is because *less_than_half* refers to the frequency of giving less than half, implying that a decline in the transfer to the poorer recipient corresponds to an increase of the frequency of *less_than_half*. In what follows, I will only state hypotheses regarding *min_giving* explicitly, though the reverse hypotheses are always implied for the outcome *less_than_half*.

As a consequence of the missing values from most individuals for at least some intervals due to the endogeneity of Δ, estimating M1 may lack sufficient power to yield significant differences with respect to the other intervals when these are considered separately. By comparing the outcome Y_{it} when $\Delta \in [0.25, 0.75)$ with the mean of all other outcomes when $\Delta \notin [0.25, 0.75)$, the problem of insufficient power is less severe. The corresponding fixed-effects model writes:

$$Y_{it} = \alpha_i + \beta_3 D_3^{it} + \epsilon_{it} \tag{M2}$$

When estimating M2, I omit all observations with $\Delta \in (-\infty, -1.75)$. The reason is that depending on an individual's aversion to inequality betwen others γ, $(-\infty, -1.75)$ is likely to overlap with $\Delta < 1 + y_i - \bar{y}$. By equation (4.9), giving to the poorer recipient may decrease and be as low 50% in this region of Δ. Thus, including these observations would mitigate the effect relative to $\Delta \notin [0.25, 0.75)$.[17] Formally, the hypothesis is:

$$\mathbf{H_0} : \beta_3 \geq 0 \ \mathbf{vs.}\ \mathbf{H_1} : \beta_3 < 0.$$

In addition, I test for rank-specific effects by interacting rank-specific dummies with the critical interval dummy D_3. Subject i's earned rank is denoted by $Rank_i$, with the time index being omitted as it remains constant over time. The model with interaction effects writes:

$$Y_{it} = \alpha_i + \sum_{r=2}^{5} \beta_r \big(D_3^{it} \times \mathbb{1}_{Rank_i = r} \big) + \epsilon_{it} \tag{M3}$$

[17]If $\Delta < -1.75$, the poorer receiver has received more that 3.75 EUR already. For this to happen within six or less preceding rounds, the richer recipient may not have received more than 2.25 EUR, putting him almost at par with the poorer recipient (net of effort).

where the coefficients β_r identify the effect of being in the critical interval $\Delta \in [0.25, 0.75)$ relative to $\Delta \notin [0.25, 0.75)$ for a particular rank r. Importantly, the base effect of one's rank is controlled for through including individual fixed effects. For the dependent variable *min_giving*, the hypotheses are:

$$\mathbf{H_0} : \beta_r \geq 0 \ \mathbf{vs.}\ \mathbf{H_1} : \beta_r < 0 \ \forall\ r = 2, 3, 4, 5.$$

To estimate the models M1, M2 and M3, I rely only on the data obtained from subjects ranked second to fifth. As the sets of recipients subjects ranked first and last must distribute to differ systematically from the choice set associated with ranks second to fifth, subjects ranked first and last do never face the inequality-rank trade-off. That is, the theoretical prediction of a local decrease in giving to the poorer recipient due to rank-loss aversion does not hold for the group of subjects ranked first or last. Hence, this group can be considered as a control group while those ranked second to second-to-last can be considered as the treatment group. To exploit the systematic treatment variation across ranks, model M4 relies on a difference-in-difference (DiD) estimation approach.

When comparing the outcomes *min_giving* and *less_than_half* across all ranks using a DiD estimator, Δ is not a useful predictor as it differs systematically for those ranked first and last by design. Yet, there exists a simple one-to-one mapping from Δ to the accumulated amount distributed to the poorer recipient. Since the difference to the poorer recipient is always equal to 2 EUR (except once), this relationship can be expressed by $G = 2 - \Delta$ for all participants ranked second to fifth, where G denotes the total amount given to the poorer recipient. Based on this relationship, the critical interval in terms of G for which theory predicts a decline in transfers to poorer recipients within the treatment group is given by $(1.25, 1.75]$.

Additionally, using the DiD approach makes it possible to identify the effects of rank-loss aversion without relying on the non-monotonicity prediction that requires to take into account decisions from all intervals. Rather, I can restrict attention to those intervals that are perfectly comparable if there was no rank-loss aversion when evauluating the behavior of the treatment group against that of the control group. The range for which I can reliably claim that the theoretical predictions would be identical if there was no rank-loss aversion is broadly given by $G \leq 1.75$. This is because for $G > 1.75$, the total earnings gap between the recipients of subjects ranked second to second-to-last has become quite small whereas it is still large between the recipients of subjects ranked first and last. Hence, other motives such as the aversion to inequality between others could drive predictions apart even if rank-loss aversion was absent in both groups.

Analogously to the analysis relying on the income difference Δ as the main predictor, the intervals used to compute *min_giving* and *less_than_half* using the DiD approach are given by $[0, 0.75]$, $(0.75, 1.25]$ and $G \in (1.25, 1.75]$. Likewise, interval dummies $\{G_j\}_{j=1,2,3}$ satisfy $G_1 = 1$ if $G \in [0, 0.75]$, $G_2 = 1$ if $G \in (0.75, 1.25]$, $G_3 = 1$ if $G \in (1.25, 1.75]$ and $G_j = 0$ otherwise. To test whether $G \in (1.25, 1.75]$ affects *min_giving* and *less_than_half* among subjects ranked second to second-to-last (treatment group), I estimate the following individual fixed-effects model:

$$Y_{it} = \alpha_i + \beta_1 \mathbb{1}_{Rank_i \in \{2,3,4,5\}} + \beta_2 \left(\mathbb{1}_{Rank_i \in \{2,3,4,5\}} \times G_3^{it} \right) + \beta_3 G_3^{it} + \epsilon_{it} \tag{M4}$$

where β_1 can be omitted due to the inclusion of individual fixed effects. The formal hypothesis is

$$\mathbf{H_0} : \beta_2 \geq 0 \textbf{ vs. } \mathbf{H_1} : \beta_2 < 0.$$

In addition, the DiD model offers an alternative test of rank-specific effects, complementing model M3. The corresponding model writes:

$$Y_{it} = \alpha_i + \sum_{r=1}^{6} \eta_r \mathbb{1}_{Rank_i = r} + \sum_{r=2}^{5} \beta_r \left(G_3^{it} \times \mathbb{1}_{Rank_i = r} \right) + \beta_3 G_3^{it} + \epsilon_{it} \tag{M4'}$$

with hypotheses identical to M3.

4.3.4 Subject Pool and Experimental Procedures

The experiment was programmed using zTree (Fischbacher, 2007). In total, 90 subjects (18 per session) participated in the experiment at the Mannheim Laboratory for Experimental Economics (mLab) from march to april, 2016. Participants were recruited via the ORSEE recruitment system (Greiner et al., 2003). Since instructions were in English, many participants were exchange students. The majority of participants were students enrolled in business or economics-related degree programs. Moreover, 50 participants were male. A session lasted 70 minutes and yielded average earnings of 11 EUR, including a show up fee of 3 EUR.

4.4 Results

I first report results on the estimation and real-effort task before exploring the effects of rank-loss aversion in the distribution task.

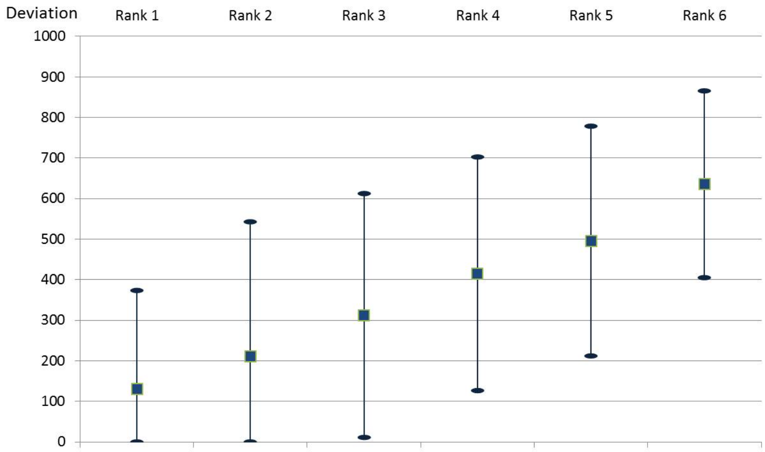

Mean deviation of guessed from true value ± 2 SEs

Figure 4.1: Accuracy in the estimation task

4.4.1 Real-Effort Task

Figure 4.1 shows that performance in the estimation task, which indirectly determines the income ranking, varies greatly even after conditioning on the rank achieved within one's group. That is, even though the estimates of subjects ranked first seem more precise on average, they are not distinctively more precise than those of the other subjects. The lack of significant differences implies that attaining a certain rank is likely to be the results of (matching) luck rather than skill. Thus, the rank assignment may be assumed to be almost random, suggesting that potential rank-specific differences are not due to self-selection.

In the real-effort task, 89 out of 90 participants accomplished to count the maximum number of tables assigned to them based on their ranking from the estimation task. The participant who failed to complete all tables still managed to count 15 out of 16 tables, leading to an income of 7.50 EUR instead of 8.00 EUR. Figure 4.2 shows the average time participants needed to complete their task. In line with the assumptions made in section 4.3.2, the data indicate an almost proportional relationship between counting time and the number of tables assigned to participants.

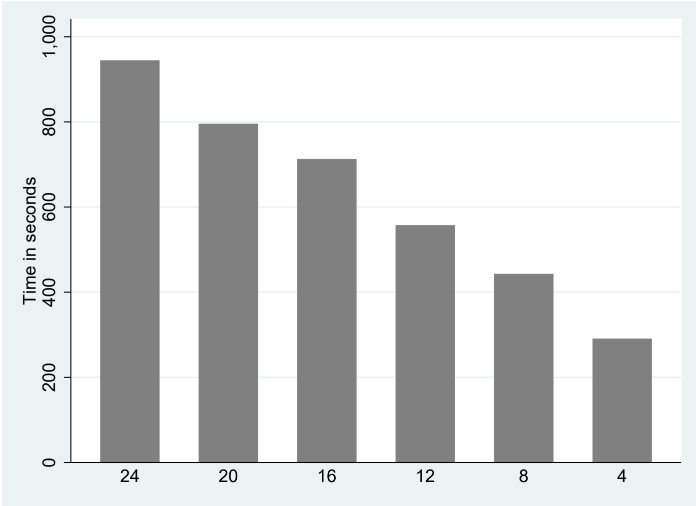

Numbers on the horizontal axis represent the assigned number of tables. The vertical axis shows the associated average total counting time.

Figure 4.2: Time in the counting task

Since all participants (with only one exception) have reached the earnings cap implied by the assigned number of tables, the ranking produced through the estimation task is identical to the pre-distribution phase income ranking. This means that the income difference between participants at adjacent ranks always amounts to 2 EUR while the effort difference always amounts to 4 tables. Thanks to this consistency, choices can easily be compared across different ranks and groups.

4.4.2 Distribution Task

The experiment asks participants to make consecutive distribution decision over multiple rounds. Since for each individual, the number of rounds equals five plus a random variable drawn from a Poisson process with mean one, only seven participants were asked to make more than seven distribution decisions. Since the trade-off between inequality aversion and rank-loss aversion always emerges in the first couple of rounds, I omit the data collected after round seven from those seven individuals in order to have a panel panel that is more balanced. Importantly, omitting these observations has no effect on differences in *min_giving* or *less_than_half* between the critical interval $[0.25, 0.75)$ and the neighboring intervals.

Figure 4.3 shows how the outcome variable *min_giving* varies with the accumulated amount (G) already given to the poorer recipient. Recall that due to the fixed earnings scheme, there is a direct link between G and the income difference to the poorer recipient $\Delta = 2 - G$ for those participants ranked second to fifth. That is, the in-

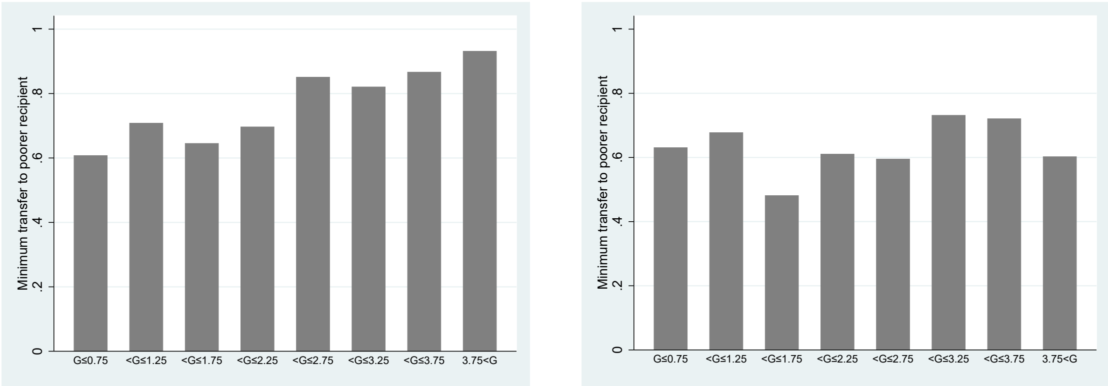

(a) Rank concerns ruled out (b) Rank concerns possible

Figure (a) contains the data from participants ranked first and sixth who never face the trade-off between inequality and rank. Figure (b) contains data from participants ranked second to fifth, who are predicted to face the rank-inequality trade-off when $G \in (1.25, 1.75] \Leftrightarrow \Delta \in [0.25, 0, 75)$. Per interval, only the lowest transfer to the poorer recipient (*min_giving*) is considered, ruling out multiple observations per interval from the same individual.

Figure 4.3: Minimum giving and accumulated transfers

terval $(1.25, 1.75]$ corresponds to the difference interval $[0.25, 0, 75)$ where the theory based on rank-loss aversion predicts that transfers to the poorer recipient from those ranked second to fifth will be lowest. According to figure 4.3b, the mean of the lowest transfers for $G \in (1.25, 1.75]$ indeed lies below that in any other interval. In contrast, no such non-monotonic pattern can be found in figure 4.3a, which contains data from participants ranked first and last. Even though looking at the raw data (figure 4.3b) suggests quite a large decline in *min_giving* of almost 20% relative to the other intervals, the empirical analysis does not consistently confirm this view. Table 4.1 contains the regression results corresponding to the identification strategy described in section 4.3.3. In column (model) M1, *min_giving* is not significantly smaller for $\Delta \in [0.25, 0.75)$ than for any other interval. Estimating model M2 is more efficient than M1 as it only assumes an effect of $\Delta \in [0.25, 0.75)$ relative to $\Delta \notin [0.25, 0.75)$, implying more observations within a single comparison category. Yet. estimating M2 confirms the results obtained from M1 as the coefficient is neither economically nor statistically significant.

Note that for estimating M1 and M2, it is implicitly assumed that the effect of rank-loss aversion is similar across ranks. Potential rank-dependent effects are analyzed in the third column of table 4.1, which presents the results obtained from estimating M3. Contrary to the research on last-place aversion by Kuziemko et al. (2014), I find no evidence that rank-loss aversion is particularly strong among subjects ranked second-to-last. Instead, I observe a significant decrease of transfers by participants ranked fourth (p-value$<$0.01, one-sided t-test). To the best of my knowledge, there is no theoretical argument for why people should be averse to giving up the fourth rank but should not be concerned about other ranks. Rather, it seems likely that the group

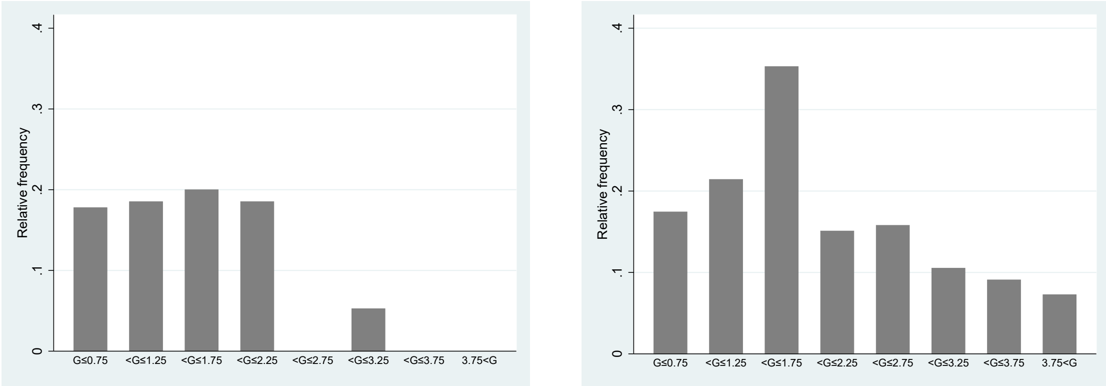

(a) Rank concerns ruled out (b) Rank concerns possible

Figure (a) contains the data from participants ranked first and sixth who never face the trade-off between inequality and rank. Figure (b) contains data from participants ranked second to fifth, who are predicted to face the rank-inequality trade-off when $G \in (1.25, 1.75] \Leftrightarrow \Delta \in [0.25, 0, 75)$. Per interval, the outcome *less_than_half* is taken into account at most once from each individual.

Figure 4.4: Frequencies of giving less than 50% to poorer recipients

of participants ranked fourth exhibits a large degree of rank loss-aversion by accident. Hence, this finding suggests that rank-loss aversion does affect the behavior of some people, even though it may not be visible on the aggregate level.

As argued in section 4.3.3, the relative frequency of choices that involve a share of less than 50% for the poorer recipient is another outcome of interest. Figure 4.4 depicts how *less_than_half* varies across different intervals for G or Δ, respectively. The large spike in *less_than_half* observed in figure 4.4b at $G \in (1.25, 1.75]$ turns out to be significant this time. Estimation results from M1 (first column) show that *less_than_half* is significantly more frequent when $\Delta \in [0.25, 0.75)$ than within most other intervals (one-sided t-test). The average magnitude of the effect may be observed most easily by focusing on the coefficient β_3 obtained from estimating M2. As displayed in the second column, *less_than_half* is chosen significantly more frequently for $\Delta \in [0.25, 0.75)$ than for $\Delta \notin [0.25, 0.75)$. With a coefficient of 0.126 (p-value<0.05, one-sided t-test) and a constant of 0.168, the result suggests that the frequency of *less_than_half* almost doubles when rank-loss aversion plays a role. Finally, the third column (M3) shows that there are almost no rank-specific effects. While the direction of the effect is positive and comparable in magnitude across all ranks, *less_than_half* increases (weakly) significantly when $\Delta \in [0.25, 0.75)$ only for participants ranked second.

As explained in section 4.3.3, I can also test for the effect of rank-loss aversion by exploiting the differences in the predictions between different ranks. Specifically, participants ranked first or sixth never need to give up their rank in order to reduce overall inequality. Hence, neither should *min_giving* decrease nor should *less_than_half*

Model	(M1)	(M2)	(M3)
Dependent variable:	min_giving	min_giving	min_giving
D_1	0.0779*		
	(0.0523)		
D_2	0.00472		
	(0.0492)		
D_4	-0.0719		
	(0.0595)		
D_5	-0.0216		
	(0.0580)		
D_6	-0.0157		
	(0.0570)		
D_7	-0.0408		
	(0.0975)		
D_8	-0.271***		
	(0.0746)		
$\mathbf{D}_3$		-0.0195	
		(0.0432)	
$D_3 \times (Rank = 2)$			0.0349
			(0.0423)
$D_3 \times (Rank = 3)$			0.0162
			(0.0897)
$D_3 \times (Rank = 4)$			-0.161***
			(0.0670)
$D_3 \times (Rank = 5)$			0.0407
			(0.107)
Constant	0.659***	0.638***	0.638***
(= D_3 for M1)	(0.0404)	(0.00514)	(0.00489)
Observations	336	294	294

Clustered SE in parentheses, One-sided test: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 4.1: Regression min_giving

Model:	(M1)	(M2)	(M3)
Dependent variable:	less_than_half (lth)	lth	lth
D_1	-0.163**		
	(0.0699)		
D_2	-0.0631		
	(0.0817)		
D_4	-0.135*		
	(0.0970)		
D_5	-0.0845		
	(0.0994)		
D_6	-0.114*		
	(0.0764)		
D_7	-0.113		
	(0.0923)		
D_8	-0.0898		
	(0.0830)		
$\mathbf{D}_3$		0.126**	
		(0.0695)	
$D_3 \times (Rank = 2)$			0.151*
			(0.109)
$D_3 \times (Rank = 3)$			0.0433
			(0.106)
$D_3 \times (Rank = 4)$			0.189
			(0.158)
$D_3 \times (Rank = 5)$			0.123
			(0.164)
Constant	0.268***	0.168***	0.168***
	(0.0648)	(0.00749)	(0.00746)
Observations	382	325	325

Clustered SE in parentheses, One-sided test: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 4.2: Regression less_than_half

increase for those subjects when $G \in (1.25, 1.75]$. This implied difference in distributive choices can easily be seen in the data by comparings the graphs depicted in figure 4.4. Moreover, table 4.3 contains the estimation results using the DiD approach M4. As shown in the first column, being in the treatment group (ranked second to fifth) causes a significant reduction by 0.2 EUR (33% when compared to the constant) in giving to the poorer recipient for $G \in (1.25, 1.75]$ (p-value<0.01, one-sided t-test). The analysis of rank-dependent effects on *min_giving* in (M4') provides further evidence in support of the previous findings. As before, the negative effect of $G \in (1.25, 1.75]$ is economically and statistically most significant for participants ranked fourth (p-value<0.01, one-sided t-test). Moreover, I find a negative effect for subjects ranked second (p-value<0.05, one-sided t-test) as well as for subjects ranked third (p-value<0.10, one-sided t-test). Perhaps surprisingly, the only group that seems consistently unconcerned about their rank giving up their rank is the group of subjects ranked second-to-last, contradicting the theory of last-place aversion. Concerning the other outcome variable *less_than_half*, the DiD method estimates a positive and significant effect of 0.201 (p-value<0.10, one-sided t-test) of $G \in (1.25, 1.75]$ ($\Leftrightarrow \Delta \in [0.25, 0.75)$), though the coefficient is less significant than those obtained from M1 and M2. In addition, estimating M4' confirms that the increase in the frequency of giving less than 50% of the money to the poorer recipient is increasing most among subjects ranked fourth and second.

To summarize, I consistently find a positive effect on the frequency of giving less than 50% due to rank-loss aversion. Yet, the negative effect on the lowest transfer is significant only by using the DiD approach or by focusing on the subset of subjects ranked fourth. Hence, rank-loss aversion seems to affect at least some people's choices, though questions remain about the magnitude of the effect on the aggregate level.

4.5 Robustness Check

The estimates obtained in the previous section strongly depend on the choice of the interval for which the effects of rank-loss aversion are predicted to be observable. Even though the respective interval has been derived based on a theory of rank-loss aversion, the discretion of the researcher when making assumptions about individual preferences challenges the credibility of the identification in section 4.4. The reason is that for any randomly generated panel, there always exists an interval for Δ in which transfers to the poorer recipient are significantly lower than the average outside of the interval if the respective interval can be chosen to be arbitrarily small. To provide some evidence that the observed effects are not just the consequence of a random process, I report re-

Model:	(M4)	(M4')	(M4)	(M4')
Dependent variable:	min_giving	min_giving	less_than_half	less_than_half
Rank $\in \{2,3,4,5\}$	0		0	
	(.)		(.)	
$\mathbf{D}_3\times$ (Rank $\in \{2,3,4,5\}$)	-0.197***		0.201*	
	(0.0727)		(0.127)	
$\mathbf{D}_3$	0.0951**	0.0951**	-1.85e-16	7.71e-17
	(0.0562)	(0.0565)	(0.105)	(0.105)
$D_3 \times (Rank = 2)$		-0.194**		0.300*
		(0.0834)		(0.187)
$D_3 \times (Rank = 3)$		-0.156*		0.101
		(0.107)		(0.124)
$D_3 \times (Rank = 4)$		-0.349***		0.288*
		(0.0957)		(0.203)
$D_3 \times (Rank = 5)$		-0.0906		0.118
		(0.121)		(0.179)
Constant	0.640***	0.640***	0.183***	0.183***
	(0.00682)	(0.00651)	(0.0111)	(0.0110)
Observations	261	261	263	263

Clustered SE in parentheses, One-sided test: *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

Table 4.3: DiD regression

sults from an alternative interval specification in this section. The alternative interval boundaries can be obtained from the previous ones by adding the constant 0.25, implying a new critical interval of $[0.5, 1)$. This interval still lies within the range of $[0, 1)$, for which theory predicts that transfers to the poorer recipient would be lower. Yet, it differs significantly from $[0.25, 0.75)$, implying that finding again an effect that is simply the consequence of a purely random process is highly unlikely. The remaining intervals and corresponding dummy variables are given by: $D_1 = 1$ if $\Delta \in [1.5, \infty)$, $D_2 = 1$ if $\Delta \in [1, 1.5)$, $D_3 = 1$ if $\Delta \in [0.5, 1)$, $D_4 = 1$ if $\Delta \in [0, 0.5)$, $D_5 = 1$ if $\Delta \in [-0.5, 0)$, $D_6 = 1$ if $\Delta \in [-1, -0.5)$, $D_7 = 1$ if $\Delta \in [-1.5, -1)$, $D_8 = 1$ if $\Delta \in [-2, -1.5]$, $D_9 = 1$ if $\Delta \in (-\infty, -2)$ and $D_j = 0 \; \forall \; j$, otherwise.

The main conclusion of this section will be that most results are robust to this alternative interval specification, suggesting that the findings do not obtain from accidentally picking the right interval. Table 4.4 shows the results from estimating the effect of $\Delta \in [0.5, 1)$ on *min_giving*. As in the previous section, it cannot be unambiguously established that transfers to the poorer recipient are significantly higher if $D_{j \neq 3} = 1$ than if $D_3 = 1$. Contrary to the main analysis however, the estimate obtained from M2 suggests that $\Delta \in [0.5, 1)$ leads to a significant decline (p<0.05, one-sided t-test) when compared to the average across all other intervals. As in table 4.1, all coefficients obtained from the rank-specific estimation M3 have a negative sign with only the coefficient for participants ranked fourth being significant (p<0.01, one-sided t-test).

Similarly, most results concerning the effect on *less_than_half* found in the main analysis can be reproduced using the alternative interval specification. According to table 4.5, participants choose *less_than_half* most often when $\Delta \in [0.5, 1)$. Except for $\Delta \in [-0.5, 0)$ and $\Delta \in [-2, 1.5)$, the obtained coefficients are always significant, including the coefficients β_2 and β_4, which refer to the intervals before and after the critical interval $[0.5, 1)$. Compared to the case $\Delta \in [1, 1.5)$, the frequency of giving *less_than_half* declines by 0.158 (p<0.05, one-sided t-test) from 0.327 for $\Delta \in [0.5, 1)$ and it declines by 0.186 (p<0.05, one-sided t-test) when compared to $\Delta \in [0, 0.5)$. Thus, the robustness check lends strong support to the non-monotonicity prediction derived in section 4.3.2. As in the main analysis, comparing $D_{j \neq 3} = 1$ with $D_3 = 1$ in (M2) yields a significant difference of 0.178 (p<0.01, one-sided t-test). Similarly, the rank-specific analysis (M3) yields positive estimates for all ranks as well. Yet, only the fourth and fifth rank are estimated to affect the outcome significantly.

Finally, the estimated coefficients of the DiD model (M4) are presented in table 4.6. Recall that the interval dummies represent the total amount transferred to the poorer recipient (G) rather than the earnings difference to the poorer recipient (Δ)

here. While this ensures the necessary comparability between participants ranked first or sixth and all other participants, it does not change the interpretation of the results due to the existing one-to-one relationship ($G = 2 - \Delta$) between G and Δ, implying that $D_3 = 1$ *iff* $G_3 = 1$ for participants ranked second to second-to-last. Notice that due to the consideration of the alternative interval specification, the condition for G corresponding to $\Delta \in [0.5, 1)$ is $G \in (1, 1.5]$. The first column shows that the effect of $G \in (1, 1.5]$ on *min_giving* is significant (p<0.05, one-sided t-test). With a magnitude of -0.177, the effect compares with the one obtained from the analogue estimation in the main analysis (-0.197). Similarly, the estimated effect on *less_than_half* also shows a significant increase by 0.354 (p<0.01, one-sided t-test). Additionally, the rank-specific analysis from estimating M4' confirms the existing findings.

Model	(M1)	(M2)	(M3)
Dependent variable:	min_giving	min_giving	min_giving
D_1	0.154***		
	(0.0504)		
D_2	0.0888**		
	(0.0433)		
D_4	0.0226		
	(0.0504)		
D_5	-0.000830		
	(0.0632)		
D_6	0.0335		
	(0.0561)		
D_7	0.0418		
	(0.0650)		
D_8	-0.131*		
	(0.0829)		
D_9	-0.209***		
	(0.0792)		
$\mathbf{D}_3$		-0.0769**	
		(0.0403)	
$D_3 \times (Rank = 2)$			-0.0489
			(0.0604)
$D_3 \times (Rank = 3)$			-0.0802
			(0.0907)
$D_3 \times (Rank = 4)$			-0.178***
			(0.0726)
$D_3 \times (Rank = 5)$			-0.00971
			(0.0839)
Constant	0.596***	0.646***	0.646***
	(0.0359)	(0.00416)	(0.00428)
Observations	330	310	294

Clustered SE in parentheses, One-sided test: *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

Table 4.4: Regression min_giving (robustness)

Model:	(M1)	(M2)	(M3)
Dependent variable:	less_than_half (lth)	lth	lth
D_1	-0.230***		
	(0.0784)		
D_2	-0.158**		
	(0.0719)		
D_4	-0.186**		
	(0.101)		
D_5	-0.115		
	(0.108)		
D_6	-0.164**		
	(0.0800)		
D_7	-0.170**		
	(0.0832)		
D_8	-0.101		
	(0.113)		
D_9	-0.161**		
	(0.0933)		
$\mathbf{D}_3$		0.178***	
		(0.0734)	
$D_3 \times (Rank = 2)$			0.113
			(0.135)
$D_3 \times (Rank = 3)$			0.174
			(0.135)
$D_3 \times (Rank = 4)$			0.271*
			(0.169)
$D_3 \times (Rank = 5)$			0.165*
			(0.126)
Constant	0.327***	0.165***	0.166***
	(0.0694)	(0.00700)	(0.00740)
Observations	382	346	324

Clustered SE in parentheses, One-sided test: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 4.5: Regression less_than_half (robustness)

Model: *Dependent variable:*	(M4) min_giving	(M4') min_giving	(M4) less_than_half	(M4') less_than_half
Rank $\in \{2,3,4,5\}$	0 (.)		0 (.)	
$\mathbf{D}_3\times$ (Rank $\in \{2,3,4,5\}$)	-0.177** (0.0825)		0.354*** (0.124)	
$\mathbf{D}_3$	0.0667 (0.0711)	0.0667 (0.0716)	-0.193** (0.108)	-0.193** (0.109)
$D_3 \times (Rank = 2)$		-0.0797 (0.0808)		0.193* (0.139)
$D_3 \times (Rank = 3)$		-0.159* (0.115)		0.296** (0.127)
$D_3 \times (Rank = 4)$		-0.330*** (0.100)		0.546*** (0.202)
$D_3 \times (Rank = 5)$		-0.140 (0.123)		0.382*** (0.138)
Constant	0.654*** (0.00684)	0.654*** (0.00654)	0.190*** (0.0104)	0.189*** (0.0103)
Observations	245	245	247	247

Clustered SE in parentheses, One-sided test: $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 4.6: DiD regression (robustness)

4.6 Conclusion

In this paper, I report results from an experiment designed to test whether people are averse to a loss of their income rank when absolute income is held constant. The design forces participants into a trade-off requiring them to forego the opportunity to decrease inequality unless they give up their rank. Even though payoffs from the experiment are not affected by a change in one's relative position, people tend to reduce the amount they distribute to the poorer recipient when distributing the normal amount would cause loss of their rank. The effect is statistically significant in some but not all regression specifications. In addition, I consistently estimate that significantly more people give less than 50% of the total transfer to the poorer of the two recipients when their rank is at stake. In other words, participants are more likely to choose an allocation that increases inequality in their group if such a choice preserves the existing ranking.

Those findings lend credibility to an entire strand of literature focusing on the implications of people competing for higher relative income. While the existing experimental evidence is limited to high status (first place) loving and low status (last place) aversion, my results suggest that people value any income rank even in the absence of any real status implication or additional material benefits. Importantly, the trade-off between inequality and rank also rules out inequality aversion as a potential explanation. Yet, the regression-dependent variations in the magnitude and significance of the effect suggest that the data is not fully conclusive and calls for further research concerning the effects of rank-loss aversion.

Appendices

A Appendix Chapter 2

A.1 Proofs of Section 2.3

Proof of Lemma 2

Proof. In **Part I**, I show that if consumers follow the stopping rule as specified in lemma 1, there is a unique sequence of increasing prices that maximize profits, denoted by $\{p_k^+\}_{k=1,\dots,N}$. Under the equilibrium stopping rule $\mathcal{R}^*$, consumers always buy if $v_{ik} > p_k$. Taking all other forthcoming sellers' expected prices as given, a price by seller k can induce a different stopping rule for at least some types $x \in X$ if $p_k > p_{k+1}^e + \delta,\ \delta > 0$ such that those types find continuing to search worthwhile despite $v_{ik} > p_k$. In **Part II**, I show that any price $p_k > p_k^+$ inducing an alternative stopping rule cannot be profitable if it is not profitable under $\mathcal{R}^*$.

<u>**Part I**</u>

Existence. Define $q_k(p) := -\frac{D_k(p)}{D_k'(p)}$ such that the FOC writes: $p = q_k(p)\ \forall\ k = 1, ..., N$. Next, write

$$q_k(p) = \frac{\int_p^1 g(x) f_k(x) \mathrm{d}x}{g(p) f_k(p)}$$

and observe that $q_k(p)$ is continuous $\forall p > 0$ and satisifes: $\lim_{p \to \underline{v}} q_k(p) > 0$ and $\lim_{p \to 1} q_k(p) = 0$. Hence, there always exists a $p \in \mathbb{R}$ that solves $p = q_k(p)$.

Uniqueness. Define $\theta_k(x) := g(x) f_k(x)\ \forall\ k \geq 1$. Then, expected demand at seller k conditional on observing a history $h = k - 1$ writes:

$$D_k(p) = \int_p^1 \theta_k(x) \mathrm{d}x = \Theta_k(1) - \Theta_k(p),$$

where $\Theta_k(\cdot)$ is the antiderivative of $\theta_k(\cdot)$. Thus, I can rewrite $q_k(p)$ as:

$$q_k(p) = \frac{\Theta_k(1) - \Theta_k(p)}{-\theta_k(p)}.$$

Notice that that $\theta_k(p)$ is log-concave for all k since first, $\theta_1(x) = f_1(x) g(x) = f(x) g(x)$ is log-concave as multiplication preserves log-concavity and because second, log-concavity of $\theta_k(p)$ implies log-concavity of $\theta_t(p)\ \forall\ t \geq k$. The second statement follows from:

$$\theta_{k+1} = g(x) f_{k+1}(x) = \frac{g(x)\big(1 - g(x)\big) f_k(x)}{\int_{\underline{v}}^1 \big(1 - g(t)\big) f_k(t) \mathrm{d}t} = \frac{1 - g(x)}{C_k} \theta_k(x),$$

where C_k is a constant. The argument in footnote 1 shows that log-concavity of $g(x)$ implies log-concavity of $1 - g(x)$ $\forall x \in \{x : g(x) \leq 1\}$. Again, multiplication of two positive and log-concave functions preserves log-concavity. Hence, log-concavity of $\theta_k(x)$ implies log-concavity of $\theta_{k+1}(x)$. Consequently, $\theta_k(p)$ is log-concave $\forall\ k \leq K^*$.

Further, log-concavity of $\theta_k(p)$ implies log-concavity of the anti-derivative $\Theta_k(p)$. Define $\Delta_k(p) = \big(\Theta_k(1) - \Theta_k(p)\big)/O_k$ where $O_k := \int_{\underline{v}}^1 g(x)f_k(x)\mathrm{d}x$ is a normalization to make $\Delta_k(p)$ a probability measure. Then, $\Delta_k(p)$ is log-concave over its positive domain[1]. Rewriting $q_k(p)$ yields:

$$q_k(p) = \frac{\Delta(p)}{-\Delta'(p)} \frac{O_k}{O_k}$$

As $\Delta_k(p)$ is a log-concave probability distribution, $q_k(p)$ is decreasing (see Bagnoli and Bergstrom (2005) and, thus, has a unique fix point.

The unique sequence of prices is increasing in k. Consider again $q_k(p)$ for arbitrary k:

$$q_k(p) \leq q_{k+1}(p) \textit{ iff } \frac{\int_p^1 g(x)f_k(x)\mathrm{d}x}{g(p)f_k(p)} \leq \frac{\int_p^1 \frac{g(x)\left(1-g(x)\right)f_k(x)}{\int_{\underline{v}}^1 \left(1-g(t)\right)f_k(t)\mathrm{d}t}\mathrm{d}x}{\frac{1-g(p)}{\int_{\underline{v}}^1 \left(1-g(t)\right)f_k(t)\mathrm{d}t} g(p)f(p)} \tag{A.1}$$

$$\Leftrightarrow \int_p^1 g(x)\big(1-g(p)\big)f_k(x)\mathrm{d}x \leq \int_p^1 g(x)\big(1-g(x)\big)f_k(x)\mathrm{d}x \tag{A.2}$$

$$\Leftrightarrow \int_p^1 g(x)\big(g(p)-g(x)\big)f_k(x)\mathrm{d}x \geq 0 \textit{ iff } g'(x) < 0 \tag{A.3}$$

Because inequality (A.3) always holds for $g'(x) < 0$, $q_k(p) \leq q_{k+1}(p)$ holds as well. The following proof goes by contradiction. Hence, assume that $p^*_{k+1} < p^*_k$ for at least some k. Then, A.3 implies that:

$$p^*_{k+1} = q_{k+1}(p^*_{k+1}) > q_k(p^*_{k+1})$$

As $q_k()$ is decreasing, it follows from the assumption of $p^*_{k+1} < p^*_k$ that:

$$q_k(p^*_{k+1}) > q_k(p^*_k)$$

[1]Subtracting $\Theta_k(1)$ preserves log-concavity. Further, multiplication with a log-concave function (i.e. (-1)) preserves log-concavity over the composed function's positive domain and O_k is just a constant.

But then uniqueness of $p = q_k(p)$ implies $p^*_{k+1} = q_{k+1}(p^*_{k+1}) > q_k(p^*_k) = p^*_k$, contradicting the assumption.

Part II

Now consider seller k setting a price $p'_k > p^+_k$ such that continuing to search is worthwhile for at least some consumers, i.e. $\exists x \in X$ s.th. $V_{k+1}(x_i) > 0$ even if $v_{ik} > p'_k$. If for some type x, $V_{k+1}(x_i) > 0$ even if $v_{ik} > p'_k$, then $x \in \hat{X}_k$, with

$$\hat{X}_k := \{x \in X : g(x)\big(p'_k - p_{k+1}\big) > s\}. \tag{A.4}$$

This is because $p^e_{k+2} > p^e_{k+1}$, and hence the surplus from sampling any seller beyond the next seller cannot be positive if the surplus from sampling only the next seller is not positive. By construction, $\exists!\ \bar{x} \in \hat{X}_k$ (with $\bar{x} > p'_k$) s.th. $x \in \hat{X} \Rightarrow x > \bar{x}$. If $x_i > \bar{x}$, consumer i might return with some probability and buy from k. Denote that probability by ρ. It is sufficient to note that $\rho \leq 1\ \forall x_i > \bar{x}$ with strict inequality for a positive mass of consumers(If $\rho = 1$ always, there would be no reason to continue search in the first place. Hence, it would not correspond to prices that lead to an alternative stopping rule for at least some consumers.) Due to $p'_k > p^e_{k+1} + \delta$, seller k's demand $\hat{D}_k(p)$ under any alternative stopping rule writes

$$\hat{D}_k(p) = \int_p^{\bar{x}} \rho \cdot g(x) f_k(x) \mathrm{d}x + \int_{\bar{x}}^{\bar{v}} g(x) f_k(x) \mathrm{d}x \tag{A.5}$$

$$< \int_p^{\bar{x}} g(x) f_k(x) \mathrm{d}x + \int_{\bar{x}}^{\bar{v}} g(x) f_k(x) \mathrm{d}x \qquad = D_k(p). \tag{A.6}$$

where $D_k(p)$ denotes the demand if consumers always followed $\mathcal{R}^*$. As shown in Part I, the problem $\max_p D_k(p)p$ has a unique solution at $p = p^+_k$ and thus:

$$p^+_k \cdot D_k(p^+_k) > p'_k \cdot D_k(p'_k)\ \forall\ p'_k \neq p^+_k$$

Therefore, it follows from inequality (A.6) that:

$$p^+_k \cdot D_k(p^+_k) > p'_k \cdot \hat{D}_k(p'_k)$$

□

Proof of Lemma 3

Proof. Define $j^* = \sup\{k : p_k^e < p_k^+, k \leq N\}$. By construction, it holds that $p_{j^*}^e < p_k^e \;\forall\; k \geq j^*$, implying that for $p_{j^*} = p_{j^*}^e$, a consumer buys from j^* if $v_{ij^*} \geq p_{j^*}^e$. Notice also that it does not matter whether a consumer with history $h = j^*-1$ has encountered other matches previously. Since $p_{j^*}^e < p_k^e \;\forall\; k > j^*$, the fact that sampling j^* must have been worthwhile to a consumer implies that $p_{j^*}^e < p_k \;\forall\; k \in \{j < j^* : v_{ij} > 0\}$. Thus, all consumers arriving at seller j^* will buy if $v_{ij^*} > p_{j^*}$ and $p_{j^*} \leq p_{j^*}^e + \delta$, where $\delta > 0$ is determined such that for prices $p_{j^*} > p_{j^*}^e + \delta$, some consumers are induced to continue searching or return to a previous seller despite $v_{ij^*} > p_{j^*}$. Hence, for any price $p_{j^*} \leq p_{j^*}^e + \delta$, seller j^* has full monopoly power. Thus, if the price belongs to this interval, it solves

$$\max_p \Big(pD_{j^*}(p) \Big)$$

Consequently, there always exists a profitable deviation from $p_{j^*} = p_{j^*}^e < p_{j^*}^e + \delta$ if $\partial_p pD_{j^*}(p)|_{p=p_{j^*}^e} > 0$. Let $\tilde{p} = \min_{k<j^*}\{p_k : v_{ik} > p_k\}$ and denote by $\hat{X} := \big\{x \in X : g(x)\big(\min\{x, \tilde{p}\} - p_{j^*}^e\big) > s\big\}$ the set of consumers whose expected surplus from sampling j^* despite available matches is positive due to potential price savings. If potential savings in the price are sufficiently low, $\hat{X} \subseteq \emptyset$. Then, consumers follow $\mathcal{R}^*$ and seller j^*'s demand is fully captured by the expression given for $D_{j^*}(p)$. From lemma 2, we know that

$$p_{j^*}^e < p_{j^*}^+ = \frac{D_{j^*}(p_{j^*}^+)}{-D'_{j^*}(p_{j^*}^+)} < \frac{D_{j^*}(p_{j^*}^e)}{-D'_{j^*}(p_{j^*}^e)} \tag{A.7}$$

where the second inequality holds as $D_{j^*}(p)$ is log-concave under $\mathcal{R}^*$ such that the RHS of the FOC is decreasing in p. Hence, $\partial_p pD_{j^*}(p)|_{p=p_{j^*}^e} > 0$, rendering $p'_{j^*} > p_{j^*}^e$ a profitable deviation.

Next, consider the case where some types' stopping behavior does change such that $\hat{X} \not\subseteq \emptyset$ and denote by $\hat{D}_{j^*}(p)$ the resulting expected demand at seller j^*. As before, denote by $f_k(x)$ the PDF of arriving consumers if they follow $\mathcal{R}^*$ such that $\hat{X} \subseteq \emptyset$ and but by $\hat{f}_k(x)$ the PDF of types if $\hat{X} \not\subseteq \emptyset$. By construction, $\exists\; \underline{x} \in \mathbb{R}$ s.th. $x \in \hat{X} \Rightarrow x \geq \underline{x}$. For $f(x)$ and $\hat{f}(x)$, this implies

$$f_{j^*}\big(x, \mathcal{R}'\big) \geq f_{j^*}\big(x, \mathcal{R}^*\big) \forall\; x \geq \underline{x},$$

with strict inequality for all $x \in \hat{X}$. By continuity of $g(x)$, $\int_{x\in\hat{X}} \mathrm{d}x > 0$ and since $\hat{f}_{j^*}(x)$ is a PDF, it follows that $\hat{f}_{j^*}(x) = z \cdot f_{j^*}(x) \;\forall\; x < \underline{x}$ where $z < 1$ is a normalization. I can thus conclude about the RHS of (A.7) that:

$$\begin{aligned}
\frac{D_{j^*}(p)}{-D'_{j^*}(p)} &= \frac{\int_p^{\underline{x}} g(x) f_{j^*}(x)\mathrm{d}x + \int_{\underline{x}}^{\bar{v}} g(x) f_{j^*}(x)\mathrm{d}x}{-g(p) f_{j^*}(x)} \\
&< \frac{\int_p^{\underline{x}} g(x) f_{j^*}(x)\mathrm{d}x + \int_{\underline{x}}^{\bar{v}} (1/z) g(x) \hat{f}_{j^*}(x)\mathrm{d}x}{-g(p) f_{j^*}(x)} \\
&= \frac{z \int_p^{\underline{x}} g(x) f_{j^*}(x)\mathrm{d}x + \int_{\underline{x}}^{\bar{v}} g(x) \hat{f}_{j^*}(x)\mathrm{d}x}{-z \cdot g(p) f_{j^*}(x)} \\
&= \frac{\hat{D}_{j^*}(p)}{-\hat{D}'_{j^*}(p)} \; \forall \; p < p^e_{j^*} + \delta
\end{aligned}$$

Combining the above inequality with (A.7) yields:

$$p^e_{j^*} < \frac{\hat{D}_{j^*}(p^e_{j^*})}{-\hat{D}'_{j^*}(p^e_{j^*})}$$

and thus $\partial_p p \hat{D}_{j^*}(p)|_{p=p^e_{j^*}} > 0$. Consequently, seller j^* would always deviate to a higher price $p > p^e_{j^*}$ if $p^e_{j^*} < p^+_{j^*}$ and hence, $p^e_{j^*} \geq p^+_{j^*}$.

As $p^e_{j^*} < p^+_{j^*}$ is eliminated, another seller $j^* = \sup\{k : p^e_k < p^+_k, k \leq N\}$ might exist. However, by going backwards and using the same argument as above, any $p^e_k < p^+_k$ can be ruled out until $j^* = 1$. □

Proof of Proposition 2.1

Proof. Begin with the first seller a consumer samples, i.e. $k = 1$. By Lemma 3, $p^e_k \geq p^+_k \; \forall \; k$ and thus $p^+_1 < p^e_k \; \forall \; k > 1$. Define $\hat{X}_1$ as in (A.4). Then $\exists! \delta > 0$ s.th. $\hat{X}_1 \subseteq \emptyset$ *iff* $p_1 \leq p^+_1 + \delta$. Note that $\hat{X}_1 \subseteq \emptyset$ induces consumers to adopt the search behavior $\mathcal{R}^*$ and by Part I of lemma 2, $p^+_1 = \text{argmax } D_1(p)p$ where $D_1(p)$ is a hypothetical demand function for $p_1 > p^+_1 + \delta$ that would prevail if continuing to search despite a match were ruled out by assumption (and consumers followed $\mathcal{R}^*$).

It remains to show that $p_1 = p^+_1$ yields larger profits than any price $p_1 > p^+_1 + \delta$, inducing a stopping rule different from $\mathcal{R}^*$ for at least some types. I use the notation of the previous proof and denote the demand that arises under an alternative stopping rule by $\hat{D}(p)$. By Part II of lemma 2, $\hat{X}_1 \not\subseteq \emptyset$ implies that $p'_1 D_1(p) > p'_1 \hat{D}_1(p'_1) \; \forall \; p'_1 > p^+_1 + \delta$ and thus

$$p^+_1 \cdot D_1(p^+_1) > p^+_1 + \delta \tag{A.8}$$

Hence, profits are maximized globally at $p_1^+ = \text{argmax } D_1(p) \cdot p$, irrespective of whether other prices could induce an alternative stopping rule.

Next, consider seller $k = 2$. Since $p_1 = p_1^+ < p_k^e \ \forall \ k > 1$, any consumer sampling $k = 2$ satisfies $v_{i1} = 0$. Hence, the distribution of arriving consumers $f_2(x)$ is equal to (2.1) as derived under the consumers' stopping rule $\mathcal{R}^*$. Also, $v_{i1} = 0$ implies that consumers never return to the previous seller and thus always buy if both $v_{i2} > p_2$ and $\hat{X}_2 \subseteq \emptyset$. Since $p_2 \leq p_2^+ + \delta \Rightarrow \hat{X}_2 \subseteq \emptyset$, the previous argument for seller $k = 1$ now applies to $k = 2$. By induction, this implies $p_k = p_k^+ \ \forall k \leq N$. □

Proof of Proposition 2.2

Proof. The FOC from the maximization problem (2.7) in the main text writes:

$$0 = \sum_{i=1}^{N} \phi_i \big(pD_i'(p) + D_i(p)\big) = p \sum_{i=1}^{N} \phi_i D_i'(p) + \sum_{i=1}^{N} \phi_i D_i(p)$$

Define $\tilde{q}(K^*, p) = \frac{\sum_{i=1}^{K^*} \phi_i D_i(p)}{-\sum_{i=1}^{K^*} \phi_i D_i'(p)}$. Then price p maximizing (2.7) must solve $\tilde{q}(K^*, p) = p$, where ϕ_k is defined by (2.6). Rewriting $\tilde{q}$ yields:

$$\tilde{q}(p) = \frac{\sum_{i=1}^{K^*} \phi_i \Theta_i(1) - \sum_{i=1}^{K^*} \phi_i \Theta_i(p)}{-\sum_{i=1}^{K^*} \phi_i \theta_i(p)}$$

As $\Theta_i(p)$ is log-concave $\forall \ i$, $\sum_{i=1}^{K^*} \phi_i \Theta_i(p)$ is log-concave. By the same argument made in proving proposition 1, this implies that $\bar{\Delta}(p) = \sum_{i=1}^{K^*} \phi_i \Theta_i(1) - \sum_{i=1}^{K^*} \phi_i \Theta_i(p)$ is log-concave. This permits to write $\tilde{q}(p) = \frac{\bar{\Delta}(p)}{-\bar{\Delta}'(p)}$. By Bagnoli and Bergstrom (2005), log-concavity of $\bar{\Delta}(p)$ is sufficient that $\tilde{q}(p)$ is decreasing in p. Further, it holds that $\lim_{p \to \underline{v}} \tilde{q}(p) > 0$ and $\lim_{p \to 1} \tilde{q}(p) = 0$. Hence, the price p solving the FOC $p = \tilde{q}(p)$ exists and is unique.[2]

□

Proof of Lemma 4

Proof. The proof is based on the following algebraic property:

$$\frac{x_1}{y_1} < \frac{x_2}{y_2} \Rightarrow \frac{x_1}{y_1} < \frac{\phi_1 x_1 + \phi_2 x_2}{\phi_1 y_1 + \phi_2 y_2} < \frac{x_2}{y_2} \quad \textit{for} \quad x, y, \phi > 0 \tag{A.9}$$

[2]The existence and unique proofs are basically identical to those used in proving lemma 2.

Suppose $\phi_1, \phi_2 > 0$. The LHS follows from the following algebra:

$$\begin{aligned}\frac{x_1}{y_1} < \frac{x_2}{y_2} \;&\Leftrightarrow \phi_1 x_1 y_2 < \phi_1 x_2 y_1\\ &\Leftrightarrow \phi_1 x_1 y_2 + \phi_2 x_1 y_1 < \phi_1 x_2 y_1 + \phi_2 x_1 y_1\\ &\Leftrightarrow x_1(\phi_1 y_1 + \phi_2 y_2) < y_1(\phi_1 x_1 + \phi_2 x_2)\\ &\Leftrightarrow \frac{x_1}{y_1} < \frac{\phi_1 x_1 + \phi_2 x_2}{\phi_1 y_1 + \phi_2 y_2}.\end{aligned}$$

The proof of the RHS of inequality (A.9) is a tautology. Iterating over inequality (A.9) implies that $\frac{\sum_{i=1}^{K^*}\phi_i D_i(p)}{-\sum_{i=1}^{K^*}\phi_i D_i'(p)} < \frac{D_{K^*}(p)}{-D_{K^*}'(p)}$. Besides, the proof of lemma 2 tells us that $\frac{D_{K^*}(p)}{-D_{K^*}'(p)} < \frac{D_{K^*+1}(p)}{-D_{K^*+1}'(p)}$. Writing $x_1 = \sum_{i=1}^{K^*}\phi_i D_i(p)$, $y_1 = -\sum_{i=1}^{K^*}\phi_i D_i'(p)$, $x_2 = D_{K^*+1}(p)$. and $y_2 = -D_{K^*+1}'(p)$, it follows from inequality (A.9) that:

$$\frac{\sum_{i=1}^{K^*}\phi_i D_i(p)}{-\sum_{i=1}^{K^*}\phi_i D_i'(p)} < \frac{\sum_{i=1}^{K^*}\phi_i D_i(p) + \phi_{K^*+1}D_{K^*+1}(p)}{-\sum_{i=1}^{K^*}\phi_i D_i'(p) + \phi_{K^*+1}D_{K^*+1}'(p)}.$$

Hence, $\tilde{q}(p, K^*) < \tilde{q}(p, K^*+1)$ and more generally, $\tilde{q}(p, K^*) < \tilde{q}(p, K_2^*)$ if $K_2^* > K_1^*$. Since profit maximization implies $\tilde{q}(p, K_1^*, \hat{h}) = p(K_1^*)$ and $\tilde{q}(p, K_2^*, \hat{h}) = p(K_2^*)$, it follows that $p(K_2^*) > p(K_1^*)$. □

Proof of Proposition 2.3

The statement follows immediately from the fact that $p(K^*)$ is increasing in K^* (lemma 4) if K^* is weakly decreasing in s. while $p(K^*)$ is increasing in K^*. The argument below, including the subsequent lemmata 13, 14 and 15, shows that K^* is weakly decreasing in s.

For the simplest case, suppose that $V_k\big(K^*, p\big(K^*(s)\big), s\big) > 0 \;\forall\; k \leq N$.[3] Then, $K^* = N$ is the unique equilibrium. Otherwise however, whether an equilibrium with active search exists, depends crucially on the ordering of the elements of $\{\hat{s}_k(K^*)\}_{k=1,..,N}$ for every K^*, where

$$\hat{s}_k(K^*) \in \{s : V_k\big(K^*, p\big(K^*(s)\big), s\big) = 0\}$$

is the threshold level of search costs specifying that for a given price $p(K^*)$ and search persistence K^*, sampling seller k is worthwhile if and only if $s \leq \hat{s}_k$.

It is instructive to begin with the case where consumers become more pessimistic while searching, meaning that $V_k\big(K^*, p\big(K^*(s)\big), s\big) < V_{k+1}\big(K^*, p\big(K^*(s)\big), s\big) \;\forall\; k < K^*$

[3] I switched back to the notion of "seller k" here since I believe it makes the analysis more tractable. Notice that this is equivalent to referring to the continuation value of a consumer with $h = k - 1$.

and for all K^*. This implies that $\hat{s}_k(K^*) > \hat{s}_{k+1}(K^*)\ \forall\ k < K^*$ such that search persistence decreases smoothly as there are no "jumps" in K greater than one.

Lemma 13 *There is a sequence of intervals* $\{(\hat{s}_{K+1}(K), \hat{s}_K(K)]\}_{K=1,..,N}$ *separated by closed neighborhoods such that consumers follow a stopping rule in pure strategies and sample up to*

$$K^* = K \text{ sellers for } s \in (\hat{s}_{K+1}(K), \hat{s}_K(K)]$$

Lemma 14 *There is a sequence of intervals* $\{(\hat{s}_K(K), \hat{s}_K(K-1)\}_{K=1,..,N}$ *disjoint from the intervals characterized in lemma 13 and separated by closed neighborhoods such that consumers follow a stopping rule in pure strategies up to seller* $K^* = K-1$ *and randomize over sampling and not sampling seller* K *with some probability* $m(K) \in (0,1)$.

Proof. Since $\hat{s}_k(K^*) > \hat{s}_{k+1}(K^*)\ \forall\ k < K^*$, $K^* < N$ only if$s > \hat{s}_N(N)$. Further, because continuation values are decreasing, $\hat{s}_{K^*}(K^*) > \hat{s}_{K^*+1}(K^*+1)\ \forall\ K^* < N$. Consequently, $K^* = N-1$ if $\hat{s}_{N-1} < s < \hat{s}_N - \epsilon,\ \epsilon > 0$.

At $s = \hat{s}_N(N)$, the the direct effect of a decrease in K^* from N to $N-1$ is zero, since the benefit of sampling seller N is exactly offset by s. However, the indirect effort through the price is strictly larger than zero. By proposition 4, it reduces the price and thus raises the continuation value such that $V_N(N-1, p(N-1), \hat{s}_N(N)) > 0$, implying that consumers would actually sample N sellers.

This inconsistency arises for all search costs in the range of $\hat{s}_N(N) \leq s < \hat{s}_N(N-1) < \hat{s}_{N-1}(N-1)$, where $\hat{s}_N(N-1)$ is determined by the general rule $\hat{s}_K(K-1) \in \{s : V_K(K, p(K-1), s) = 0\}$.

For $s \in \big[\hat{s}_N(N), \hat{s}_N(N-1)\big)$, consumers thus choose a mixed stopping rule, sampling seller N only with some probability $m(N) \in (0,1)$. From substituting $\phi'_N = \phi_N \cdot m(N)$ for ϕ_N into equation (2.7) and looking at the resulting FOC, it follows that the uniform optimal price $p(m(N))$ is decreasing in $m(N)$ and always satisfies $p(m) \in \big(p(N-1), p(N)\big)\ \forall\ m \in (0,1)$.

Since $V_N\big(N, p\big(N\big), s\big) < 0 < V_N\big(N, p\big(N-1\big), s\big)$ for $s \in \big[\hat{s}_N(N), \hat{s}_N(N-1)\big)$, there thus always exists an $m(N) \in (0,1)$ such that $V_N\big(N, p(m), s\big) = 0$ and $V_k\big(N, p(m), s\big) > 0\ \forall\ k < N$.

Applying a randomized stopping rule prior to any seller $k < N$ cannot be an equilibrium strategy. Since $V_k\big(N, p(N), s\big) > 0\ \forall\ k < N$, it follows that $V_k\big(K^*, p(m(k)), s\big) > 0$ for any $K^* < N$ and $m(k) < 1$ for some seller k by proposition 4. However, randomization at k requires $V_k = 0$, thus leading to a contradiction.

By the same argument, consumers use a only pure stopping strategies if $s \in (\hat{s}_N(N-1), \hat{s}_{N-1}(N-1)]$ and (conditional on $v_{ik} = 0 \; \forall \; k < N-1$) randomize over sampling an not sampling seller $N-1$ for $s \in (\hat{s}_{N-1}(N-1), \hat{s}_{N-1}(N-2)]$. If the threshold levels are ordered, this patterns repeats until $K^* = 1$. □

In general, the sequence of threshold levels $\{\hat{s}_k(K^*)\}_{k=1,..,N}$ need not be decreasing for every K^*. For every possible search persistence K', define $\underline{\hat{s}}(K') \in \min_k \left\{\hat{s}_k(K')\right\}_{k=1,..,K'}$. If $\hat{s}_{K'}(K') = \underline{\hat{s}}(K')$, the stopping rule is given by lemma 4 and 5 for all s such that $K^*(s) \geq K'$ with $K^*(s) \in \mathbb{K}$ as defined in (2.5). However, if there exists a $j < K'$ with $\hat{s}_j(K') = \underline{\hat{s}}(K^*)$, consumers do not begin to randomize over sampling seller K' if $s > \hat{s}_{K'}(K')$ as specified in lemma 5. Instead, already for $\hat{s}_{K'}(K')s > s_j(K')$, they randomize over sampling seller j and then continue sampling sellers up to K'. Formally, denote the upper bound on search persistence in the latter case by $\underline{K}^*(s) \in \max_K \big\{K : \hat{s}_K(K) > \underline{\hat{s}}(K)\big\}$ and denote by $j \in \{k : \hat{s}_k(\underline{K}^*) = \underline{\hat{s}}(\underline{K}^*\}$ the seller with the lowest threshold given $\underline{K}^*$. Lemma 15 summarizes the general randomized stopping rule:

Lemma 15 *Consumers follow a stopping rule in pure strategies up to seller $j-1$ and randomize over continuing to sample seller j with some probability $m(j) \in (0,1)$ for $s \in (\hat{s}_j(\underline{K}^*), \hat{s}_j(j-1)]$. Consumers who sample j continue sampling sellers up to $k = \underline{K}^*$ if $v_{ik} = 0 \; \forall \; k < \underline{K}^*$.*

Proof. Consumers with history $h = \underline{K}^*$ would find sampling seller $\underline{K}^*$ worthwhile if $s \leq \hat{s}_{\underline{K}^*}(\underline{K}^*)$. However, consumers do not "reach" seller $\underline{K}^*$ when the price is $p(\underline{K}^*)$ since $\hat{s}_j(\underline{K}^*) < \hat{s}_{\underline{K}^*}(\underline{K}^*)$ by construction. That is, they would stop sampling at seller $j < \underline{K}^*$.

For $j > 1$, the issue is resolved with a unique mixed strategy. Randomizing over the decision to sample seller j also affects demand at all seller $k' > j$. While by construction $V_k\big(\underline{K}^*, p\big(m(j), \underline{K}^*\big), s\big) > 0 \; \forall \; j < k \leq \underline{K}^*$ is unfeasible since consumers are not indifferent, the mass of consumers is reduced by the same fraction $1 - m(j)$ for all sellers $k > j$.

As in lemma 14, there exists a $m(j) \in (0,1)$ such that $V_j\big(\underline{K}^*, p(m(j)), s\big) = 0$ for $s > \hat{s}_j(\underline{K}^*$ and where $p(m)$ maximizes sellers profits. Notice that for $m(j) \to 0$, $K^* = j-1$ effectively. Hence, the threshold search level for which no mixed strategy $m(j) \in (0,1)$ can yield $V_j\big(\underline{K}^*, p(m(j)), s\big) = 0$ is given by $\hat{s}_j(j-1)$ since$V_j\big(j-1, p(j-1), \hat{s}_j(j-1)\big) = 0$ by construction. For $s < \hat{s}_j(j-1)$, consumers follow again a pure stopping strategy with $K^* = j-1$ by lemma 13.

The mixed strategy equilibrium is unique. To see why, note that generally, $V_{j'} \neq V_j$ for $j \neq j'$ and that randomizing over the decision to sample any seller j' requires that $V_{j'} = 0$ in equilibrium. (Note that I drop some arguments of V here).

Suppose that randomizing over sampling $j' > j$ was an equilibrium and that $m(j') < 1$. Then, a price ensuring $V_{j'}\big(p(m(j'))\big) = 0$ implies $V_j(p(m(j')) < 0$ since by construction of j, $V_j(p) < V_k(p)\ \forall\ k \leq \underline{K}^*$ for any price p.[4] Hence, consumers would neither sample seller j nor any $j' > j$, which is a contradiction to $m(k) \in (0,1)$. Next, suppose that consumers randomize at some $k' < j$. Then, $V_{k'}\big(p(m(k'))\big) = 0$ implies $V_j\big(p(m(k'))\big) < 0$. However, then consumers sample at most $k' < j$ sellers, and it follows from proposition 4 that $p(m(k')) < p(m(j))\ \forall\ m(j) \in [0,1]$. But $V_j\big(p(m(k'))\big) > 0$, which is a contradiction. □

To conclude, consumers' search persistence is always weakly decreasing in the level of search costs s, though it may involve mixing over sampling additional sellers or discontinuous jumps if search costs are above a certain threshold. Hence, proposition 2.3 obtains.

Proof of Proposition 2.4

Proof. Note that $p(m(j))$ is increasing in $m(j)$ only if $j > 1$. Too see why, consider the FOC for a randomized stopping rule $m(j)$:

$$p = \frac{\sum_{i=1}^{j-1} \phi_i D_i(p) + m(j) \sum_{i=j}^{K^*} \phi_i D_i(p)}{-\Big(\sum_{i=1}^{j-1} \phi_i D_i'(p) + m(j) \sum_{i=j}^{K^*} \phi_i D_i'(p) \Big)} \tag{A.10}$$

If $j = 1$, $m(j)$ cancels out from the RHS of (A.10). Thus, randomizing with $m(j) < 1$ has no effect on on the equilibrium price. Consequently, no mixed strategy $m(j)$ for $j = 1$ exists that renders $V_1 = 0$. Moreover, as shown in the proof of lemma 13 and 14, there also exists no mixed strategy for $m(k) \in (0,1)$ for $k > j = 1$. □

[4] Seller j is defined in the main text.

A.2 Proofs of Section 2.4

Proof of Proposition 2.5

Proof. Using inequality (A.9), it follows immediately that:

$$\frac{D_{1^*}(p)}{-D'_{1^*}(p)} < \frac{\sum_{i=1}^{K^*} \phi_i D_i(p)}{-\sum_{i=1}^{K^*} \phi_i D'_i(p)} < \frac{D_{K^*}(p)}{-D'_{K^*}(p)}$$

Since by lemma 2, prices with and without tracking are unique, the stated result obtains. □

Proof of Lemma 5

Proof. The proof is provided in the main text. Suppose that consumers followed the stopping rule $\mathcal{R}^*$, irrespective of the sellers price, thereby eliminating price competition entirely as in the unique equilibrium with tracking. In principle, sellers can choose any price, including the price they would choose if tracking were not available. However, even when assuming that consumers follow $\mathcal{R}^*$, sellers choose different prices to maximize profits by proposition 2.5. Hence, they cannot obtain less profits if K^* under search with tracking is at least as high as K^* under search without tracking.

□

Proof of Proposition 2.6

Proof. The result is an immediate consequence of lemma 2.4. If given the same level of search costs, the market is active under tracking while it is inactive under no tracking, strictly more matches are realized under tracking, thus leading to both higher consumer surplus and higher profits. □

Formulas used in Surplus Computations

Consumer surplus is captured by the continuation value prior to sampling the first sellers: Substituting $k = 1$ in equation (2.4) und using the specification for $f()$ and $g()$ as provided in the main text, I obtain

$$V_1 = \sum_{k=0}^{K^*} \left(\int_{p_k}^{1} (1-x)(x-p_k)\big(x\big)^k \mathrm{d}x - s\int_0^1 \big(x\big)^k \mathrm{d}x \right).$$

Under search with tracking, a seller in position k (observing $h = k - 1$) maximizes the following profit function:

$$\pi_k = p_k \cdot \left(\int_{p_k}^{1} (1-x)(x)^{k-1} \frac{1}{\int_0^1 (x)^{k-1} \mathrm{d}t} \mathrm{d}x \right). \tag{A.11}$$

To derive the uniform price without tracking conditional on search persistence K^*, the above profit functions from equation must be weighted by ϕ_k. I do not normalize conditional on K^* because the sellers FOC would remain unchanged and because I look for industry profit, not for seller profit per consumer.

$$\Pi(K^*, p) = p \cdot \left(\sum_{k=1}^{K^*} \phi_k \cdot \int_{p}^{1} (1-x)(x)^{k-1} \frac{1}{\int_0^1 (x)^{k-1} \mathrm{d}t} \mathrm{d}x \right) \tag{A.12}$$

where the probability ϕ_k is given in its general form in (2.6) and now writes $\phi_k = \int_0^1 (x)^{k-1} \mathrm{d}x$. Hence, overall profits without tracking take a very simple form:

$$\Pi(K^*, p) = p \cdot \left(\sum_{k=1}^{K^*} \cdot \int_{p}^{1} (1-x)(x)^{k-1} \mathrm{d}x \right) \tag{A.13}$$

The formulas above are implemented in a *Mathematica* code to derive optimal prices for every possible K^*. Using expressions for V_k, the equilibrium K^* is computed for every level of search costs. The code can be obtained from the corresponding author upon request.

A.3 Proofs of Section 2.5

Proof of Lemma 6

Proof. **(a) Existence and uniqueness** If consumers follow $\mathcal{R}^*$, seller k's demand writes:

$$D_k(p) = \int_{p}^{\bar{v}} g(x) \frac{\big(1-g(x)\big)^{k-1} f(x)}{\int_{\underline{v}}^{1} \big(1-g(t)\big)^{k-1} f(t) \mathrm{d}t} \mathrm{d}x$$

Profit maximization with respect to the above demand function yields first-order conditions which are equivalent to those shown in the proof of the existence of an increasing price sequence in proposition 1. Hence, a solution to

$$\max_p D_k(p) p$$

exists. Further, since the FOC $\frac{D_k(p)}{-D_k'(p)}$ is decreasing, it is also unique. Moreover, the same reasoning as in Part II regarding the possibility of setting to a price that changes the stopping rule applies. Consequently, uniqueness is preserved when accounting for deviations from $\mathcal{R}^*$.

(b) Decreasing price sequence

If $g'(x) > 0$, the inequality in (A.3) is reversed. Hence,

$$q_j(p) > q_{j+1}(p) \ \forall \ p \in [\underline{v}, \bar{v}].$$

Since the solution to $q_k(p) = p$ is unique $\forall \ j$, it follows that $p_k < p_{k-1} \ \forall \ k \leq N$ in equilibrium. □

Proof of Lemma 7

Proof. By lemma 6, $\{p_k^*\}_{k=1,\ldots,N}$ is optimal if consumers follow $\mathcal{R}^*$ always. As in any PBE, expectations are correct, it suffices to show that given sellers' optimal prices if $\hat{\Delta} > g'(x) \ \forall \ x \in X$, applying $\mathcal{R}^*$ is optimal for consumers.

Using previous notation $\hat{X}_k = \{x \in X : g(x)\big(p_k' - p_{k+1}\big) > s\}$, recall that

$$\hat{X} \subseteq \emptyset \ \Leftrightarrow \ \mathcal{R}^* \textit{ is optimal.}$$

Hence, all types apply $\mathcal{R}^*$ if

$$g(x)\big(p_k^e - p_{k+1}^e\big) \leq s \ \forall \ k < N, \ x \in \ X$$

As $g'(x) > 0$, the type $x_i = \bar{v}$ has the highest probability of encountering a match. Hence, $\hat{X} \subseteq \emptyset$ if $g(\bar{v})\big(p_k^e - p_{k+1}^e\big) \leq s \ \forall \ k$. Equation (A.3) shows that the difference in prices is a function of the slope of $g(x)$. In particular, $|p_k^e - p_{k+1}^e \,| \to 0 \textit{ if } \ g'(x) \to 0 \ \forall \ x$. Hence, it is possible to find a $g(x)$ sufficiently flat such that $\hat{X} \subseteq \emptyset$ for every $s > 0$. □

Proof of Lemma 8

Proof. Replace p_k^+ by p_k^* in the proof of lemma 3. The result follows immediately. The reason why the same argument as in the proof of lemma 3 applies is given in the main text. □

Proof of Proposition 2.6

Proof. Replace p_k^+ by p_k^* in the proof of Proposition 2. The result follows immediately. The reason why the same argument as in the proof of Proposition 2 applies is given in the main text. □

A.4 Proofs of Section 2.6

Proof of Lemma 11

Proof. **(a) Existence and Uniqueness:** Random consumer search from the perspective of sellers for all histories $h \geq \hat{h}$ implies that a seller at position $k > \hat{h}$ computes the probability that his position is k conditional on $K^* \geq k > \hat{h}$. This probability is denoted by $\phi_k(\hat{h})$ with:

$$\phi_k(\hat{h}) = \phi_k / \sum_{i=\hat{h}+1}^{K^*} \phi_i \ \forall \ k > \hat{h} \tag{A.14}$$

where ϕ_i is defined as in (2.6). Notice that as $\phi_k < \phi_{k+1}$ by construction, it also holds that $\phi_k(\hat{h}) < \phi_{k+1}(\hat{h})$. Conditional on not observing the search history h, the optimal price solves the following FOC:

$$\frac{\sum_{i=\hat{h}+1}^{K^*} \phi_i(\hat{h})\Theta_i(1) - \sum_{i=\hat{h}+1}^{K^*} \phi_i(\hat{h})\Theta_i(p)}{-\sum_{i=\hat{h}+1}^{K^*} \phi_i(\hat{h})\theta_i(p)} = p.$$

Analogously to the previous arguments, it follows that the optimal price exists and is unique for every $\hat{h}$.

(b) Comparison between $p(\hat{h})$ and $p_{\hat{h}+1}$: $p(\hat{h})$ is uniquely defined by:

$$p(\hat{h}) = \frac{\sum_{i=\hat{h}+1}^{K^*} \phi_i D_i\big(p(\hat{h})\big)}{-\sum_{i=\hat{h}+1}^{K^*} \phi_i D_i'\big(p(\hat{h})\big)}.$$

Using $D_{\hat{h}+1}(p)/-D'_{\hat{h}+1}(p) < ... < D_{K^*}(p)/-D'_{K^*}(p)$ and applying inequality (A.9) again implies:

$$\frac{\sum_{i=\hat{h}+1}^{K^*} \phi_i D_i(p)}{-\sum_{i=\hat{h}+1}^{K^*} \phi_i D_i'(p)} > \frac{D_{\hat{h}+1}(p)}{-D'_{\hat{h}+1}(p)} \ \forall \ p \in X$$

As $p_{\hat{h}+1} = \frac{D_{\hat{h}+1}(p_{\hat{h}+1})}{-D'_{\hat{h}+1}(p_{\hat{h}+1})}$, it follows by the same argument as in proposition 1 that $p(\hat{h}) > p_{\hat{h}+1} \ \forall \ 1 \leq \hat{h} < K^* - 1$. If $\hat{h} = K^* - 1$, the FOC determining $p(\hat{h})$ reduces to

$p = \frac{D_{K^*}(p)}{-D'_{K^*}(p)}$ and is thus identical to the FOC for seller K^* if the history $\hat{h}$ is disclosed. Hence, $p(\hat{h}) = p_{K^*}$ for $\hat{h} = K^* - 1$. □

Proof of Proposition 2.10

Proof. Suppose that $\hat{h} < K^*$. In equilibrium, the strategy $\hat{h}$ must be optimal, i.e. $p^e(NT) < p^e_{\hat{h}+1}$. In any symmetric equilibrium, sellers' beliefs must satisfy $\mu(h) = 0 \;\forall\; h < \hat{h}$ and $\mu(h) = \phi_{h+1} \;\forall h \geq \hat{h}$, where ϕ_{h+1} equals a seller's probability of being in position $h+1$ in a consumer's search process as defined in (2.6). Then by Lemma 7, the seller's optimal price conditional on $\hat{h}$ satisfies $p(NT) = p(\hat{h}) \geq p_{\hat{h}+1} \;\forall\; \hat{h}$. If $\hat{h} = K^*$, $p(\hat{h}) \geq p_{\hat{h}+1}$ does not affect a consumer's choice because she does not sample seller $\hat{h}+1$. However if $\hat{h} < K^*$, the optimal no tracking price $p(NT)$ contradicts the expectation of $p^e(NT) < p^e_{\hat{h}+1}$, which is necessary to sustain the equilibrium strategy $\hat{h} < K^*$. Besides, a price $p(NT) > p_{K^*}$, despite being an action that is never chosen in equilibrium, cannot be part of the seller's equilibrium strategy. If a seller observes the off-equilibrium choice NT, the maximum possible history can be $h = K^* - 1$ as otherwise the consumer would have ended search. Consequently, setting $p'(NT) = p_{K^*} < p(NT)$ constitutes a profitable deviation. Notice also that unique full disclosure outcome is robust to assuming the opposite tie-breaking rule in favor of no disclosure if $p^e(NT) = p_{\hat{h}+1}$. Changing the tie-breaking rule in that way allows for an equilibrium with $\hat{h} = K^* - 1$. However, consumers with a history $h = \hat{h}$ are the only consumers choosing no disclosure. Hence, the choice of $d = NT$ perfectly reveals the type, sellers set $p(NT) = p_{\hat{h}+1} = p_{K^*}$ and the outcome is equivalent to the unique equilibrium under the alternative tie-breaking rule.

□

Proof of Lemma 12 and Proposition 2.11

Reversing the inequality $h \geq \hat{h}$ in the proofs of Lemma 11 and Proposition 2.10 (to $h \leq \check{h}$) yields the results.

B Appendix Chapter 3

B.1 Formal Analysis of Bayesian Beliefs

For simplicity, assume that the score s is a continuous variable and denote a participant's realized score by $\hat{s}$. Further, assume that the distribution of s in the population is given by $F(s)$ with the density function $f(s)$. Denote by $\bar{m}$ the median score derived from this distribution.

Proof of Observation 1

Proof. Denote by $\mathbb{P}(A)$ denote the unconditional probability of some event A, where $A \in \{E, R\}$ represents either the earnings condition E or R. Let $\mu(E|\hat{s} \ \wedge \ H)$ be the posterior that the true state of the world (session) is E conditional on having received a high income (H). Conversely, denote by $\mu(E|\hat{s} \ \wedge \ L)$ the respective posterior belief conditional on a low income (L).

Then the posterior conditional on H writes:

$$\mu(E \mid \hat{s} \ \wedge \ H) = \frac{\mathbb{P}(\hat{s} \ \wedge \ H \mid E)}{\mathbb{P}(\hat{s} \ \wedge \ H)} = \frac{\int^{\hat{s}} f(s)\mathrm{d}s}{\mathbb{P}(E) \cdot \int^{\hat{s}} f(s)\mathrm{d}s + \mathbb{P}(R) \cdot \frac{1}{2}},$$

where $\int^{\hat{s}} f(s)\mathrm{d}s$ equals the probability of having a higher score than a randomly drawn opponent, which equals the probability of receiving a high income in E.

Notice that $\mathbb{P}(R) = \mathbb{P}(E) = \frac{1}{2}$ and consider the following two possibilities for illustration. If $\hat{s}$ is above the median score $\bar{m}$, then $\int^{\hat{s}} f(s)\mathrm{d}s > \frac{1}{2}$ and thus $\mu(E \mid \hat{s} \wedge H) > \frac{1}{2}$. Conversely, $\hat{s} < \bar{m}$ implies $\mu(E \mid \hat{s} \ \wedge \ H) < \frac{1}{2}$. □

Proof of Observation 2

Proof. A low income spectator's Bayesian posterior is given by

$$\mu(E \mid \hat{s} \ \wedge \ L) = \frac{\mathbb{P}(\hat{s} \ \wedge \ L \mid E)}{\mathbb{P}(\hat{s} \ \wedge \ L)} = \frac{\int_{\hat{s}} f(s)\mathrm{d}s}{\mathbb{P}(E) \cdot \int_{\hat{s}} f(s)\mathrm{d}s + \mathbb{P}(R) \cdot \frac{1}{2}}.$$

Hence, a low income and $\hat{s} > \bar{m}$ imply that $\mu(E \mid \hat{s} \wedge L) < \frac{1}{2}$ whereas $\mu(E \mid \hat{s} \wedge L) > \frac{1}{2}$ follows if $\hat{s} < \bar{m}$. □

B.2 Paper and On Screen Instructions

Allgemeine Instruktionen

Wir begrüen Sie herzlich und danken Ihnen für Ihre Teilnahme. Ihr Verdienst bei diesem Experiment wird von Ihren Entscheidungen abhängen. Bitte folgen Sie daher den Instruktionen aufmerksam. Wir zahlen Sie am Ende der Session privat und in bar aus. Bitte sprechen Sie während des Experiments nicht mit den anderen Teilnehmern. Wenn Sie zu einem beliebigen Zeitpunkt Fragen haben sollten, melden Sie sich bitte.

Experiment

Im Experiment werden Sie in Gruppen von je zwei Teilnehmern eingeteilt, sodass jede Gruppe aus einer Person A und einer Person B besteht. Das Experiment setzt sich aus zwei Teilen zusammen. Der erste Teil beinhaltet eine Produktionsphase, in der Sie Geld verdienen können. Im zweiten Teil können Sie das verdiente Geld innerhalb einer anderen Gruppe zwischen Person A und Person B umverteilen. Danach ist das Experiment beendet; es gibt also nur einen Durchgang.

Erster Teil

Im ersten Teil besteht Ihre Aufgabe darin, fiktive Wörter verschlüsseln, indem Sie die einzelnen Buchstaben der Wörter durch dreistellige Nummern ersetzen. Welche Nummer für welchen Buchstaben steht, wird Ihnen in einer unter dem Wort stehenden Tabelle angezeigt. Zu Beginn gibt es eine zweiminütige Probephase, in der Sie Zeit haben, sich an die Aufgabe zu gewöhnen.

Beispiel: Ihnen wird das Wort LAF angezeigt. Entsprechend der auf Ihrem Bildschirm angezeigten Tabelle gilt: L=418, A=109, F=215. Um die Lösung einzugeben, tragen Sie diese Nummern unter die Buchstaben in die dafür vorgesehenen Felder ein und klicken auf Weiter.

Sobald Sie ein Wort korrekt verschlüsselt haben, gibt Ihnen das Programm ein neues Wort zum Verschlüsseln vor. Auerdem ändert sich die Tabelle, die Ihnen sagt, mit welcher Nummer welcher Buchstabe übersetzt werden kann. Dabei werden sowohl die Reihenfolge der Buchstaben, welche nicht alphabetisch sein muss, als auch die zugehörigen Nummern variiert. Sie müssen also bei jedem Wort aufs Neue nachsehen, durch welche Nummer die gesuchten Buchstaben zu übersetzen sind. Wenn Sie auf Weiter klicken und einen Fehler bei der Verschlüsselung gemacht haben, werden Ihre Eingaben gelöscht. Sie müssen dasselbe Wort dann erneut verschlüsseln. Die Verschlüsselungstabelle ändert sich in diesem Fall nicht.

Sie erhalten solange neue Wörter zum Verschlüsseln, bis die Zeit von **25 Minuten** (1500 Sekunden) abgelaufen ist. **Für jedes korrekt verschlüsselte Wort erhalten Sie einen Punkt.**

Einkommen aus der Produktionsphase: Ihr Verdienst aus der Produktionsphase wird durch Ihre relative Punktzahl innerhalb Ihrer Gruppe oder vom Zufall bestimmt. Ob Zufall über Ihren Verdienst entscheidet, hängt davon ab, ob Sie sich in Welt 1 oder Welt 2 befinden. In Welt 1 hängt Ihr Verdienst (und auch der Verdienst aller anderen Teilnehmer) ausschlielich von der relativen Punktzahl innerhalb der Gruppe ab. Die Person mit der höheren Punktzahl erhält €20 und die Person mit der niedrigeren Punktzahl erhält 5. In Welt 2 wird zufällig entschieden, welche Person €20 und welche Person €5 erhält. Alle Teilnehmer haben dieselbe Chance, die €20 zu erhalten, unabhängig davon, ob sie Person A oder B sind. Die erreichte Punktzahl spielt keine Rolle bei der Bestimmung des Einkommens.

Die Auswahl von Welt 1 oder Welt 2 ist für alle Teilnehmer dieses Durchgangs identisch, sodass **sich entweder alle Teilnehmer in Welt 1 oder alle Teilnehmer in Welt 2 befinden.**

Beispiel (Die Zahlen sind absichtlich unrealistisch gewählt): Nehmen Sie an, dass in einer Gruppe Person A 1000 Punkte und Person B 2000 Punkte erreicht. In Welt 1 würde Person B sicher €20 und Person A sicher €5 erhalten. In Welt 2 kann es hingegen passieren, dass Person A €20 und Person B €5 erhält. Genauso wahrscheinlich ist es jedoch, dass Person B €20 erhält und Person A €5.

Ob Sie sich in diesem Durchgang des Experiments in Welt 1 oder Welt 2 befinden, erfahren Sie nicht. Diese Entscheidung ist im Vorfeld getroffen worden, wobei die Auswahl von Welt 1 oder Welt 2 gleichwahrscheinlich war.

Nach Abschluss der Produktionsphase sehen Sie lediglich Ihre eigene Punktzahl und Ihren Verdienst, nicht jedoch die Punktzahl der anderen Person in Ihrer Gruppe. Sollten Sie sich in Welt 2 befinden, erfahren Sie auch nicht, ob die-/derjenige mit der geringeren oder die-/derjenige mit der höheren Punktzahl das Einkommen von €20 erhalten hat.

Zweiter Teil

Im zweiten Teil können Sie, wenn Sie möchten, Einkommen innerhalb einer **anderen Gruppe** umverteilen. Dabei ist es ausgeschlossen, dass Mitglieder aus der Gruppe, für die Sie eine Umverteilungsentscheidung treffen, über die Umverteilung in Ihrer eigenen Gruppe entscheiden.

Über die Teilnehmer dieser anderen Gruppe erfahren Sie lediglich deren Verdienst aus der Produktionsphase, nicht jedoch die erreichten Punktzahlen.

Bei der Umverteilungsentscheidung legen Sie über einen Schiebebalken am unteren Bildschirmrand fest, ob und wieviel Geld von der Person, die das Einkommen von €20 erzielt hat, an die Person mit dem Einkommen von €5 transferiert werden soll. Der Schiebebalken erlaubt Ihnen, den Transfer in 10-Cent-Schritten zu bestimmen. Sie können den Balken entweder durch Anklicken und Gedrückthalten der linken Maustaste, oder durch Klicken der rechts und links davon angeordneten Pfeile bewegen. Sie können maximal €15 umverteilen.

Für jede Gruppe gibt es eine Umverteilungsentscheidung von insgesamt 2 Teilnehmern. Das heit, dass für die Gruppe, für die Sie einen Transfer festgelegt haben, noch die Entscheidung eines anderen Teilnehmers existiert. Der Computer bestimmt dann zufällig, wessen Entscheidung umgesetzt wird. Da entweder Ihre Entscheidung oder die des anderen Teilnehmers umgesetzt wird, können Sie ihre Transferentscheidung unabhängig davon treffen, was Sie glauben, wie dieser andere Teilnehmer entscheiden wird. Auerdem bedeutet das, dass für Ihre Wahl eine 50%ige Chance besteht, tatsächlich umgesetzt zu werden. Bitte entschieden Sie daher mit Sorgfalt.

Wenn Sie zum aktuellen Zeitpunkt Fragen haben, heben Sie bitte jetzt die Hand. Anderenfalls bitten wir Sie, das Experiment zu starten und zunächst die Verständnisfragen zu beantworten. Ziel der Fragen ist es sicherzustellen, dass Sie die verschiedenen Komponenten des Experiments korrekt verstanden haben. Ihre Entscheidungen und Antworten während des Experiments bleiben anonym.

Weitere Hinweise

- Während der Produktionsphase wird Ihnen die verbleibende Zeit in Sekunden am oberen rechten Bildschirmrand angezeigt
- Sie können das Nummernfeld Ihrer Tastatur verwenden. Dieses ist bereits aktiviert.
- Sie können die Tabulator-Taste verwenden, um bei der Eingabe der Nummern zum nächsten Eingabefeld zu gelangen

Figure B.1: Encryption task

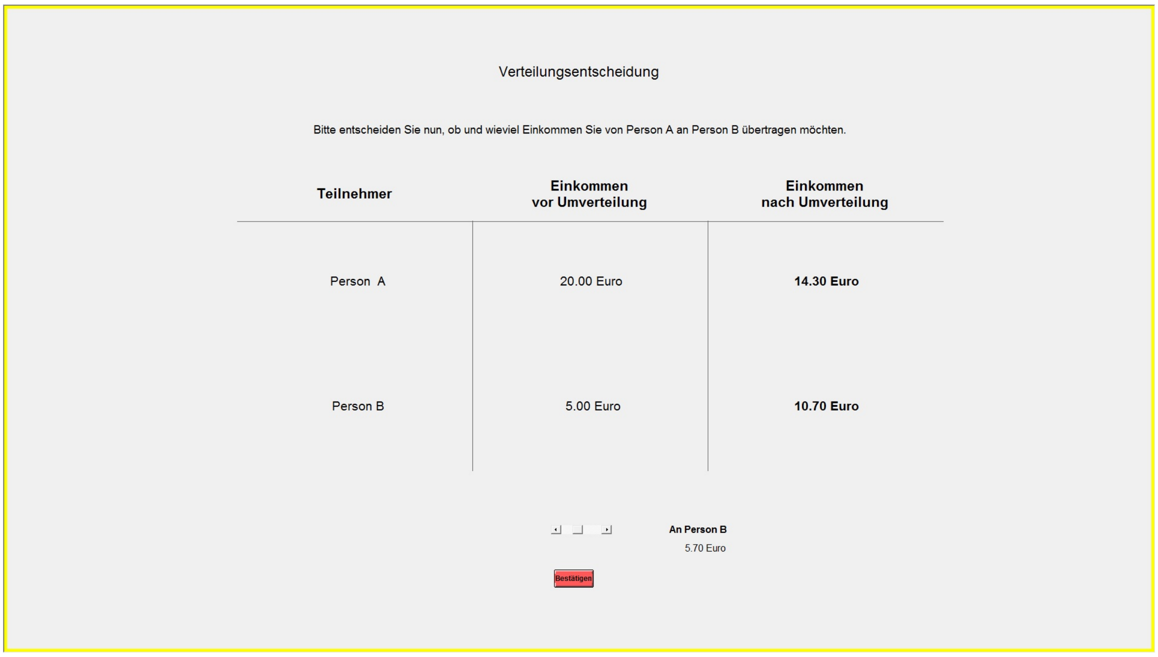

Figure B.2: Redistribution stage

Figure B.3: Belief elicitation I

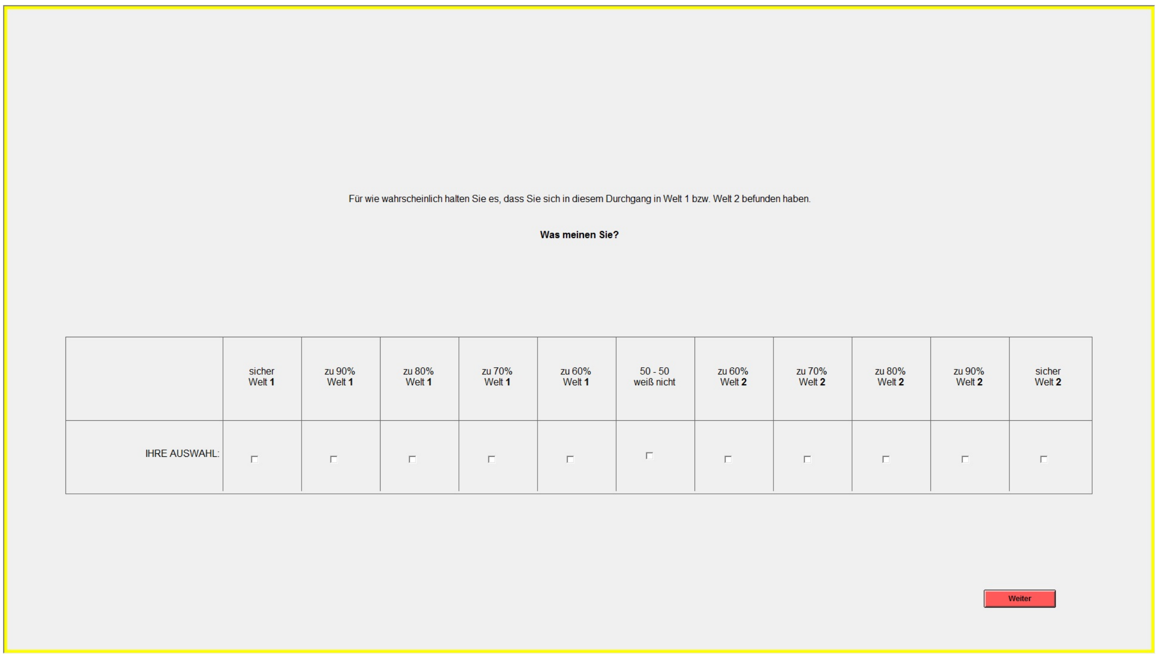

Figure B.4: Belief elicitation II

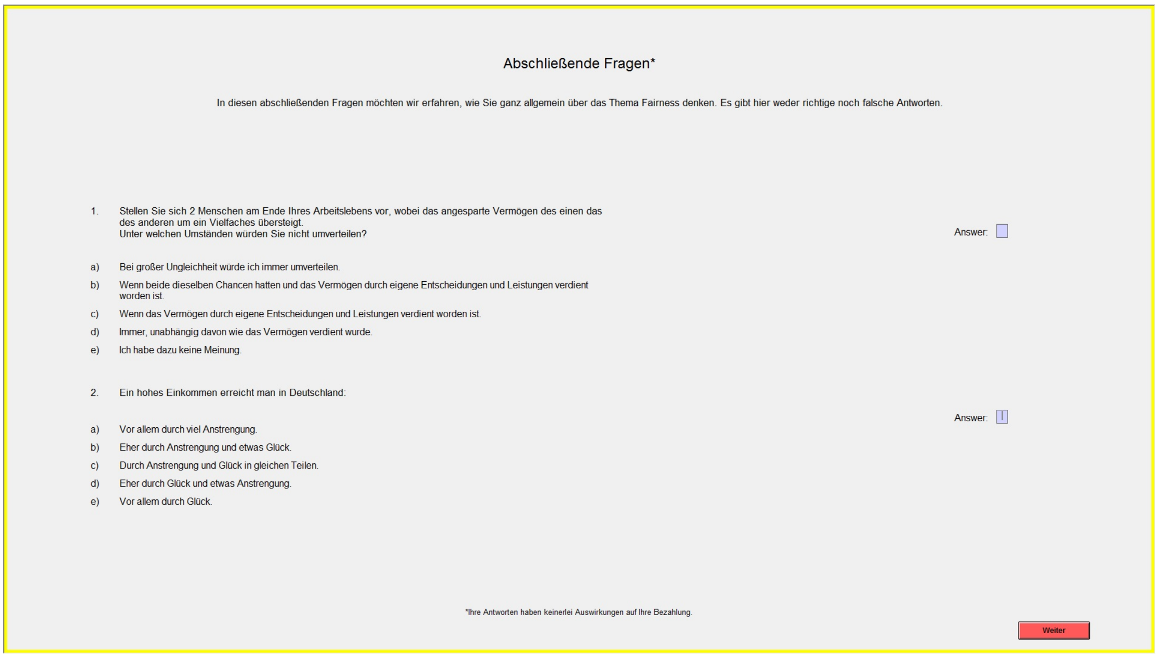

Figure B.5: Survey questions

C Appendix Chapter 4

C.1 Paper and On Screen Instructions

General Instructions

Thank you for agreeing to take part in this study. You will receive €3 as a show-up fee for participating in this session. You may also receive additional money, depending on the decisions you and other participants make. Upon completion of the session, this additional amount and the show-up fee will be paid to you individually and in private.

A clear understanding of the following instructions will help you make better decisions and increase your earnings. Please do not speak to other participants during the experiment.

Part I

The experiment is divided into two parts. Below you will find the instructions for the first part, consisting of Stage 1, Stage 2, and Stage 3. Once the first part is completed, you will receive instructions for the second part (Stage 4). You will participate in each stage only once. At the beginning of the experiment, the computer will randomly match you with five other people in the room. That is, you will be part of a group of six people. To identify each other while remaining anonymous, each one of you will be assigned a unique name from the following list: Yellow, Blue, Green, Red, Orange, and Brown.

Stage 1 - Estimation Task

Your relative performance in the estimation task within your group determines how much you can work and earn during stage 2. Details on the earnings during stage 2 are given in the description of stage 2.

The task in stage 1 is the same for all participants. You will see an identical table consisting of zeros and ones for **15 seconds**. After the time has passed, the table will disappear and you will be asked to give an estimate of how many zeros there were in the table. Based on your and the other group members estimates, a ranking will be built that assigns the first rank to the group member whose estimate has been closest to the true value, the second rank to the group member whose estimate has been the second closest to the true value, etc. In the unlikely event of a tie, the computer allocates the higher rank based on a random process with equal chances for everyone.

Example 1: (Note that the numbers may be very different in the actual stage 1.) Assume that the players Green and Red both estimate that there are 8 zeros in the table while the other players estimates are Yellow: 6, Brown: 5, Blue: 4, and Orange: 3. Furthermore, assume that the true number of zeros is 9. Then either Player Red or Player Green will be ranked first with a 50% probability while the other one will be ranked second. Player Yellow will be ranked third, player Brown fourth, player Blue fifth and player Orange sixth.

Stage 2 - Counting Task

In this part of the experiment your task is to count zeros in a series of tables. The tables are different from the one you will have seen in stage 1. The screenshot below shows the work screen you will see later (numbers in the actual stage 2 may be very different):

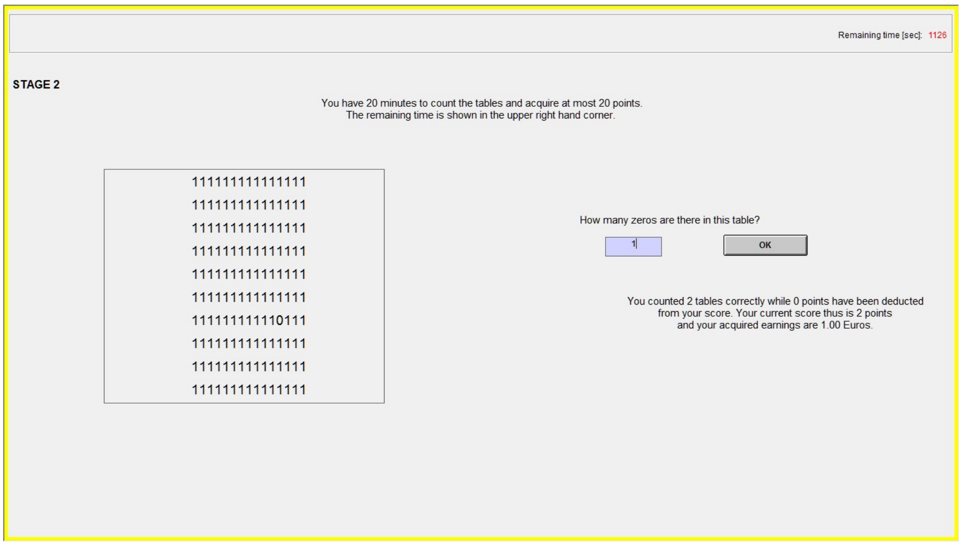

After counting the zeros, you enter the number of zeros into the box on the right hand side of the screen. After you have entered the number, you click the OK button. If you enter the correct result, a new table will be generated. If your input was wrong, you have two additional attempts to enter the correct number. You therefore have a total of three attempts to solve each table. If you enter the correct number of zeros, you earn **1 point which is worth 50 cents**. If you enter three times a wrong number for a table, 1 point (=50 cents) will be **subtracted** from your earnings and a new table will then be generated.

How many tables you can count in stage 2 (Maximum Score) depends on the rank you acquired in stage 1:

Rank	Maximum Score	Maximum possible earnings
1st place	24	12€
2nd place	20	10€
3rd place	16	8€
4th place	12	6€
5th place	8	4€
6th place	4	2€

For example, if your stage 1 rank has been 3rd you will be asked to count zeros in 16 tables. And because each correct answer is worth 50 cents you can therefore earn up to 16 50 = 800 cents = 8 €. In the unlikely event that you end up with a negative point score at the end of stage 2, your earnings will be set to 0 €. That is, you cannot accumulate negative income.

The Maximum Score includes points that are subtracted from your earnings. That is, if your Maximum Score is 16 and you miscount two tables three times, you will be able to see up to 18 tables in total before you reach 16 points (= 8 €) and finish stage 2. Independently of your rank, you will have 20 minutes (= 1200 seconds) to complete stage 2. The remaining time to complete this stage will be displayed in the upper right hand corner of the screen.

We encourage you to complete all the tables given to you. But if you do not want to complete them all, you can simply stop at any time. If you decide to stop (or if you finish early) please just wait and remain silent as you cannot leave this stage before the 20 minutes are up or everybody completed stage 2. To make waiting a little less boring for those who complete their maximum number of tables early we give you access to an article about Mannheim taken from the English Wikipedia website.

Stage 3 - Distribution Game

At the beginning of Stage 3, you will be informed of how many points and how much money each participant in your group has acquired up to that point. During Stage 3 you will participate in the “distribution game”. The distribution game will be played for a certain number of rounds, and this number is determined randomly. You will learn about the number of rounds only after the last round has been completed. Hence, in every round you should make your choices as if this round was the last round.

During each round you will be presented with a choice about how to distribute an amount of money between two other members of your group. This money is additional money and is not taken away from the amount of money you have received. All the money has to be distributed between the two and therefore your own earnings will not be affected by your choice no matter what you do.

In each round, the amount of money you will be asked to distribute and the people between whom you have to split the amount will remain the same. However, **the money you distribute to them accumulates over the rounds**.

The figure below shows a typical screenshot from stage 3 (numbers in the actual stage 3 may be very different):

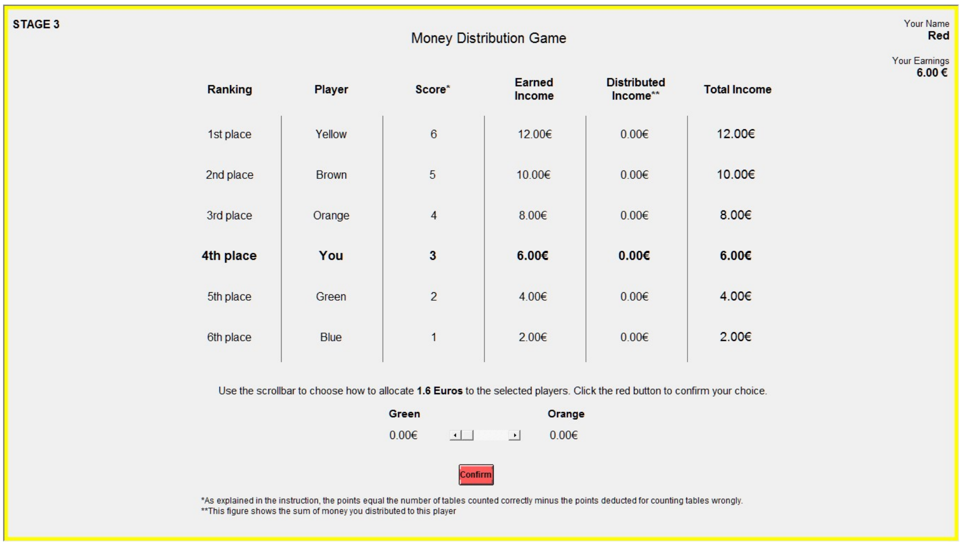

To choose how to split the money, use the scrollbar at the bottom of the screen. By moving the bar to the left, you can increase the amount distributed to the player whose name appears on the left hand side and vice versa. Before you click on the scrollbar for the first time, the computer screen will display €0 under both names, which means that no money has been distributed yet (see screenshot). However, you will not be able to click on “Confirm” before you use the scrollbar at least once in order to choose how to split the money. To make fine adjustments, you may use the left/right arrow buttons next to the scrollbar. While using the scrollbar, the columns “Distributed Incom” and “Total Income” will update automatically. In order to implement any choices you made, click on “Confirm”.

After everyone in your group has made his/her choices (including the last round), the computer will randomly select **one players choices** and award the additional money according to that players decisions in all rounds. That is, only one players distribution decisions will be implemented for the entire group. This means that whenever your decisions are chosen by the computer, everyones total income within your group will be exactly as you chose while the other players decisions will be disregarded.

However, neither you nor the other players in your group will be informed of the computers random draw at that point. This information will be made available only at the end of the experiment.

General Remarks

1. Note that all decisions will remain anonymous.
2. Always use a "." (dot) to enter decimal numbers.
3. You may earn or lose additional money in stage 4 with some probability. However, the significantly larger share of your earnings is determined during the first three stages of the experiment.

You now have another 4 minutes to read the instructions again on your own. If you have any questions now or during the experiment, please just raise your hand and I will come to your desk to answer them in private.

Figure C.1: Estimation task

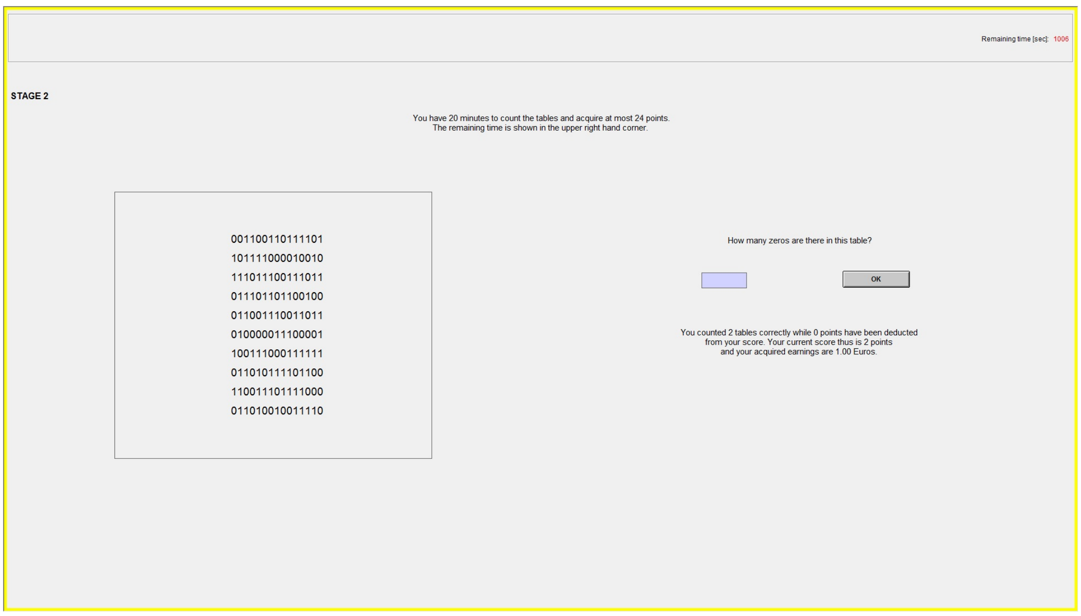

Figure C.2: Counting task

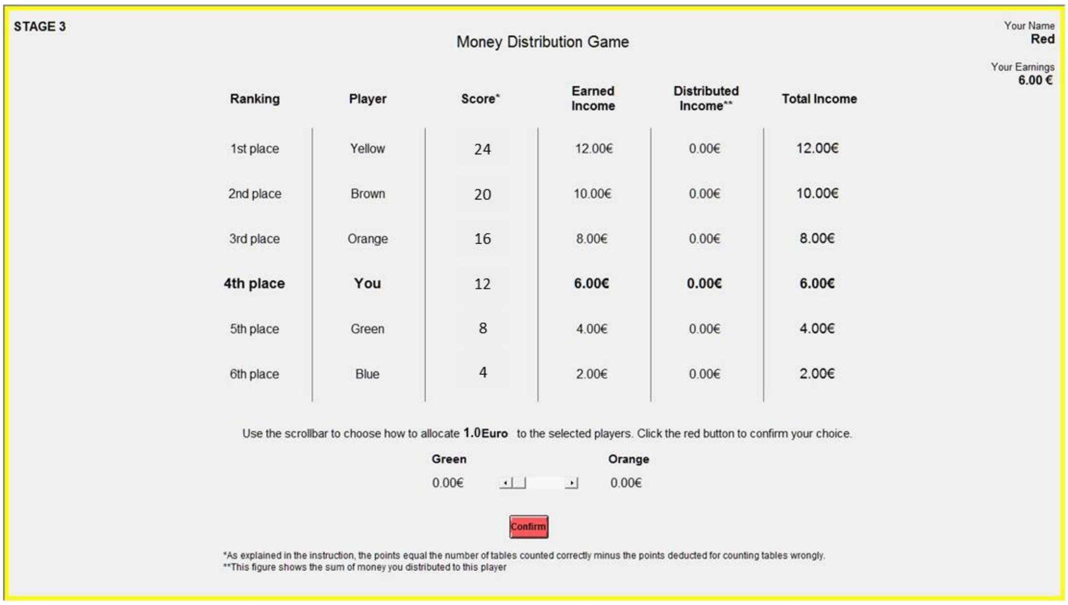

Figure C.3: Distribution stage

Bibliography

Abeler, J., Falk, A., Goette, L., and Huffman, D. (2011). Reference points and effort provision. *The American Economic Review*, pages 470–492.

Acquisti, A. and Varian, H. R. (2005). Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381.

Aguiar, M. and Hurst, E. (2007). Life-cycle prices and production. *American Economic Review*, 97(5):1533–1559.

Alesina, A. and Angeletos, G.-M. (2005). Fairness and redistribution. *The American Economic Review*, 95(4):960–980.

Alesina, A. and La Ferrara, E. (2005). Preferences for redistribution in the land of opportunities. *Journal of public Economics*, 89(5):897–931.

Almås, I., Cappelen, A. W., Sørensen, E. Ø., and Tungodden, B. (2010). Fairness and the development of inequality acceptance. *Science*, 328(5982):1176–1178.

Anderson, S. P. and Renault, R. (1999). Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics*, pages 719–735.

Anderson, S. P. and Renault, R. (2015). Search direction. Technical report, mimeo.

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

Arbatskaya, M. (2007). Ordered search. *The RAND Journal of Economics*, 38(1):119–126.

Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association*, 15(5):989–1024.

Armstrong, M., Vickers, J., and Zhou, J. (2009). Prominence and consumer search. *The RAND Journal of Economics*, 40(2):209–233.

Armstrong, M. and Zhou, J. (2010). Conditioning prices on search behaviour. Technical report, University Library of Munich, Germany.

Armstrong, M. and Zhou, J. (2011). Paying for prominence. *The Economic Journal*, 121(556).

Auten, G. E., Clotfelter, C. T., and Schmalbeck, R. L. (2000). Taxes and philanthropy among the wealthy. *Does atlas shrug*, pages 392–424.

Bagnoli, M. and Bergstrom, T. (2005). Log-concave probability and its applications. *Economic theory*, 26(2):445–469.

Ball, S., Eckel, C., Grossman, P. J., and Zame, W. (2001). Status in markets. *The Quarterly Journal of Economics*, 116(1):161–188.

Ball, S. and Eckel, C. C. (1998). The economic value of status. *The Journal of socio-economics*, 27(4):495–497.

Ball, S. B. and Eckel, C. C. (1996). Buying status: Experimental evidence on status in negotiation. *Psychology & Marketing*, 13(4):381–405.

Bar-Isaac, H., Caruana, G., and Cuñat, V. (2012). Search, design, and market structure. *The American Economic Review*, 102(2):1140–1160.

Belleflamme, P. and Peitz, M. (2015). *Industrial organization: markets and strategies*. Cambridge University Press.

Belleflamme, P. and Vergote, W. (2016). Monopoly price discrimination and privacy: The hidden cost of hiding. *Economics Letters*, 149:141–144.

Benndorf, V., Rau, H. A., and Sölch, C. (2014). Minimizing learning behavior in experiments with repeated real-effort tasks.

Blanchflower, D. G. and Oswald, A. J. (2004). Well-being over time in britain and the usa. *Journal of public economics*, 88(7):1359–1386.

Bolton, G. E. and Ockenfels, A. (2000). Erc: A theory of equity, reciprocity, and competition. *American economic review*, pages 166–193.

Borgesius, F. Z. and Poort, J. (2017). Online price discrimination and eu data privacy law. *Journal of Consumer Policy*, 40(3):347–366.

Boyce, C. J., Brown, G. D., and Moore, S. C. (2010). Money and happiness rank of income, not income, affects life satisfaction. *Psychological Science*, 21(4):471–475.

Bradley, G. W. (1978). Self-serving biases in the attribution process: A reexamination of the fact or fiction question. *Journal of personality and social psychology*, 36(1):56.

Buckley, E. and Croson, R. (2006). Income and wealth heterogeneity in the voluntary provision of linear public goods. *Journal of Public Economics*, 90(4):935–955.

Cappelen, A. W., Hole, A. D., Sørensen, E. Ø., and Tungodden, B. (2007). The pluralism of fairness ideals: An experimental approach. *The American Economic Review*, 97(3):818–827.

Cappelen, A. W., Konow, J., Sørensen, E. Ø., and Tungodden, B. (2013). Just luck: An experimental study of risk-taking and fairness. *The American Economic Review*, 103(4):1398–1413.

Cappelen, A. W., Sørensen, E. Ø., and Tungodden, B. (2010). Responsibility for what? fairness and individual responsibility. *European Economic Review*, 54(3):429–441.

Card, D., Mas, A., Moretti, E., and Saez, E. (2012). Inequality at work: The effect of peer salaries on job satisfaction. *American Economic Review*, 102(6):2981–3003.

Cassar, L. and Klein, A. H. (2017). A matter of perspective: How experience shapes preferences for redistribution. *Working paper*.

Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly journal of Economics*, pages 817–869.

Chen, Y. and He, C. (2011). Paid placement: Advertising and search on the internet. *The Economic Journal*, 121(556).

Chen, Y. and Li, S. X. (2009). Group identity and social preferences. *The American Economic Review*, 99(1):431–457.

Chen, Y. and Zhang, Z. J. (2009). Dynamic targeted pricing with strategic consumers. *International Journal of Industrial Organization*, 27(1):43–50.

Clark, A. E. and Oswald, A. J. (1996). Satisfaction and comparison income. *Journal of public economics*, 61(3):359–381.

Conitzer, V., Taylor, C. R., and Wagman, L. (2012). Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science*, 31(2):277–292.

Corneo, G. and Gruner, H. P. (2000). Social limits to redistribution. *American Economic Review*, 90(5):1491–1507.

Corneo, G. and Grüner, H. P. (2002). Individual preferences for political redistribution. *Journal of public Economics*, 83(1):83–107.

De Cornière, A. (2016). Search advertising. *American Economic Journal: Microeconomics*, 8(3):156–88.

Deffains, B., Espinosa, R., and Thöni, C. (2016). Political self-serving bias and redistribution. *Journal of Public Economics*, 134:67–74.

Diamond, P. A. (1971). A model of price adjustment. *Journal of economic theory*, 3(2):156–168.

Dijk, O., Holmen, M., and Kirchler, M. (2014). Rank matters–the impact of social competition on portfolio choice. *European Economic Review*, 66:97–110.

Durante, R., Putterman, L., and Weele, J. (2014). Preferences for redistribution and perception of fairness: An experimental study. *Journal of the European Economic Association*, 12(4):1059–1086.

Easterlin, R. A. (1974). Does economic growth improve the human lot? some empirical evidence. In *Nations and households in economic growth*, pages 89–125. Elsevier.

Eckel, C. C., De Oliveira, A., and Grossman, P. J. (2007). Is more information always better? an experimental study of charitable giving and hurricane katrina. *Southern Economic Journal*, 74(2).

Ellison, G. and Ellison, S. F. (2014). Match quality, search, and the internet market for used books1.

Engelmann, D. and Strobel, M. (2004). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *American economic review*, pages 857–869.

Erkal, N., Gangadharan, L., and Nikiforakis, N. (2011). Relative earnings and giving in a real-effort experiment. *The American Economic Review*, pages 3330–3348.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly journal of Economics*, pages 817–868.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 10(2):171–178.

Fisman, R., Jakiela, P., Kariv, S., and Markovits, D. (2015). The distributional preferences of an elite. *Science*, 349(6254):aab0096.

Fong, C. (2001). Social preferences, self-interest, and the demand for redistribution. *Journal of Public economics*, 82(2):225–246.

Frank, R. H. (1985). The demand for unobservable and other nonpositional goods. *The American Economic Review*, 75(1):101–116.

Fudenberg, D. and Tirole, J. (2000). Customer poaching and brand switching. *RAND Journal of Economics*, pages 634–657.

Greiner, B. et al. (2003). *The online recruitment system ORSEE: a guide for the organization of experiments in economics*. Max-Planck-Inst. for Research into Economic Systems, Strategic Interaction Group.

Grund, C. and Sliwka, D. (2005). Envy and compassion in tournaments. *Journal of Economics & Management Strategy*, 14(1):187–207.

Guinote, A., Cotzia, I., Sandhu, S., and Siwa, P. (2015). Social status modulates prosocial behavior and egalitarianism in preschool children and adults. *Proceedings of the National Academy of Sciences*, 112(3):731–736.

Haan, M. A. and Moraga-González, J. L. (2011). Advertising for attention in a consumer search model. *The Economic Journal*, 121(552):552–579.

Hart, O. D. and Tirole, J. (1988). Contract renegotiation and coasian dynamics. *The Review of Economic Studies*, 55(4):509–540.

Heffetz, O. and Frank, R. H. (2011). Preferences for status: Evidence and economic implications. In *Handbook of social economics*, volume 1, pages 69–91. Elsevier.

Hopkins, E. and Kornienko, T. (2004). Running to keep in the same place: Consumer choice as a game of status. *American Economic Review*, 94(4):1085–1107.

Johnson, J. P. and Myatt, D. P. (2006). On the simple economics of advertising, marketing, and product design. *The American Economic Review*, 96(3):756–784.

Kirchler, M., Lindner, F., and Weitzel, U. (2018). Rankings and risk-taking in the finance industry. *Journal of Finance*.

Konow, J. (2000). Fair shares: Accountability and cognitive dissonance in allocation decisions. *The American Economic Review*, 90(4):1072–1091.

Konow, J. (2008). The moral high ground: an experimental study of spectator impartiality.

Kuziemko, I., Buell, R. W., Reich, T., and Norton, M. I. (2014). last-place aversion: Evidence and redistributive implications. *The Quarterly Journal of Economics*, 129(1):105–149.

Larson, N. (2013). Niche products, generic products, and consumer search. *Economic Theory*, pages 1–40.

Luttmer, E. F. (2005). Neighbors as negatives: Relative earnings and well-being. *The Quarterly journal of economics*, 120(3):963–1002.

Mikians, J., Gyarmati, L., Erramilli, V., and Laoutaris, N. (2012). Detecting price and search discrimination on the internet. In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*, pages 79–84. ACM.

Milgrom, P. and Roberts, J. (1986). Relying on the information of interested parties. *The RAND Journal of Economics*, pages 18–32.

Miller, D. T., Ross, M., et al. (1975). Self-serving biases in the attribution of causality: Fact or fiction. *Psychological bulletin*, 82(2):213–225.

Mollerstrom, J., Reme, B.-A., and Sørensen, E. Ø. (2015). Luck, choice and responsibilityan experimental study of fairness views. *Journal of Public Economics*, 131:33–40.

Montes, R., Wilfried, S.-Z., and Valletti, T. (2017). The value of personal information in markets with endogenous privacy.

Moraga-González, J. L. and Petrikaitė, V. (2013). Search costs, demand-side economies, and the incentives to merge under bertrand competition. *The RAND Journal of Economics*, 44(3):391–424.

Moraga-González, J. L., Sándor, Z., and Wildenbeest, M. R. (2017). Prices and heterogeneous search costs. *The RAND Journal of Economics*, 48(1):125–146.

Mujcic, R. and Frijters, P. (2012). Economic choices and status: measuring preferences for income rank. *Oxford Economic Papers*, 65(1):47–73.

Piff, P. K., Kraus, M. W., Côté, S., Cheng, B. H., and Keltner, D. (2010). Having less, giving more: the influence of social class on prosocial behavior. *Journal of personality and social psychology*, 99(5):771.

Piff, P. K., Stancato, D. M., Côté, S., Mendoza-Denton, R., and Keltner, D. (2012). Higher social class predicts increased unethical behavior. *Proceedings of the National Academy of Sciences*, 109(11):4086–4091.

Piketty, T. (1995). Social mobility and redistributive politics. *The Quarterly journal of economics*, pages 551–584.

Rablen, M. D. and Oswald, A. J. (2008). Mortality and immortality: The nobel prize as an experiment into the effect of status upon longevity. *Journal of Health Economics*, 27(6):1462–1471.

Rawls, J. (1971). A theory of justice, harvard. *Press, Cambridge.*

Robson, A. J. (1992). Status, the distribution of wealth, private and social attitudes to risk. *Econometrica: Journal of the Econometric Society*, pages 837–857.

Robson, A. J. (2001). The biological basis of economic behavior. *Journal of Economic Literature*, 39(1):11–33.

Smith, A. (1759). The theory of moral sentiments (glasgow edition), dd raphael and a. macfie.

Taylor, C. and Wagman, L. (2014). Consumer privacy in oligopolistic markets: Winners, losers, and welfare. *International Journal of Industrial Organization*, 34:80–84.

Taylor, C. R. (2004). Consumer privacy and the market for customer information. *RAND Journal of Economics*, pages 631–650.

Turow, J., King, J., Hoofnagle, C. J., Bleakley, A., and Hennessy, M. (2009). Americans reject tailored advertising and three activities that enable it.

Veblen, T. (1899). The theory of the leisure class. *History of Economic Thought Books*.

Villas-Boas, J. M. (1999). Dynamic competition with customer recognition. *The Rand Journal of Economics*, pages 604–631.

Wolinsky, A. (1986). True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics*, pages 493–511.

Yang, H. (2013). Targeted search and the long tail effect. *The RAND Journal of Economics*, 44(4):733–756.

Zhou, J. (2011). Ordered search in differentiated markets. *International Journal of Industrial Organization*, 29(2):253–262.

Erklärung

Ich versichere hiermit, dass ich die Dissertation selbstständig und ohne Benutzung anderer als der angegebenen Quellen und Hilfsmittel angefertigt und die den benutzten Quellen wörtlich oder inhaltlich entnommenen Stellen als solche kenntlich gemacht habe. Diese Arbeit hat in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegen.

Mannheim, den 25.04.2018

Curriculum Vitae

Marcel Preuß

Since 2017	Research Assistant Prof. Dr. Dietmar Fehr University of Heidelberg, Germany
Since 2016	Teaching Assistant Chair of Economic Policy University of Mannheim, Germany
Since 2015	Research Assistant SFB 884 Political Economy of Reforms University of Mannheim, Germany
Since 2014	Doctoral Student in Economics under supervision of Prof. Dr. Henrik Orzen and Prof. Dr. Martin Peitz at the Center for Doctoral Studies in Economics (CDSE) Graduate School of Economic and Social Sciences (GESS) University of Mannheim, Germany
2012 – 2014	Master in Economic Research (M.Sc.) University of Mannheim, Germany
2011	Visiting Student Drexel-University, United Stated
2009 – 2012	Bachelor in Business Administration (B.Sc.) University of Mannheim, Germany
2009	Abitur Alexander-von-Humboldt Gymnasium, Eberswalde