

What Makes a Price Fair? An Experimental Study of Transaction Experience and Endogenous Fairness Views*

Holger Herz and Dmitry Taubinsky

October 11, 2016

Abstract

People’s fairness preferences are an important constraint for what constitutes an acceptable economic transaction, yet little is known about how these preferences are formed. In this paper, we provide clean evidence that previous transactions play an important role in shaping perceptions of fairness. Buyers used to high market prices, for example, are more likely to perceive high prices as fair than buyers used to low market prices. Similarly, employees used to high wages are more likely to perceive low wages as unfair. Our data further allows us to decompose this history dependence into the effects of pure observation vs. the experience of payoff-relevant outcomes. We propose two classes of models of path-dependent fairness preferences—either based on endogenous fairness reference points or based on shifts in salience—that can account for our data. Structural estimates of both types of models imply a substantial deviation from existing history-independent models of fairness. Our results have implications for price discrimination, labor markets, and dynamic pricing.

Keywords: Reference Points, Fairness, Salience, Bargaining, Endogenous Preferences, Price Stickiness

JEL Classification Codes: D, C9, C78

*Holger Herz: University of Fribourg and CESifo. E-mail: holger.herz@unifr.ch. Dmitry Taubinsky: University of California Berkeley and NBER. E-mail: dtaubinsky@gmail.com. We thank Philippe Aghion, Bjoern Bartling, Tom Cunningham, Stefano DellaVigna, Ernst Fehr, Oliver Hart, Michael Kosfeld, David Laibson, Michael Luca, Alex Peysakhovich, Al Roth, Klaus Schmidt, Florian Zimmermann, as well as seminar participants at Berkeley, Frankfurt, Fribourg, Harvard, Mainz, Vienna and Zurich, as well as participants at IMEBE 2013, THEEM 2013, the ESA World meeting 2013, the ESEM World meeting 2015, as well as at the V IIBEO workshop for helpful comments and discussions. We thank the Center for Foundations of Economic Preferences at the University of Zurich as well as the Harvard Law School Program on Negotiation and the Harvard Business School doctoral program for financial support of the project. Herz acknowledges support from the ERC grant on the Foundations of Economic Preferences. A previous version of this paper was circulated under the title “Market Experience is a Reference Point in Judgments of Fairness.”

1 Introduction

A large body of evidence shows that people’s aversion toward unfair transactions can play an important role in markets and negotiations. In product markets, consumers’ feelings of entitlement restrict sellers’ ability to exploit changes in supply and demand (Kahneman et al., 1986), while in labor markets, reciprocal gift exchange can lead to involuntary unemployment (Akerlof, 1982; Fehr et al., 1993).¹ To incorporate such non-pecuniary concerns into economic theory, economists have proposed models of “social preferences,” which assume that in addition to maximizing consumption, people also care about the fairness or kindness of own or others’ actions. A common property of these models is that the fairness or kindness of an action or outcome is evaluated by an exogenous and static criterion such as equal division or surplus maximization, which implies that fairness judgments remain stable over time and past experiences should not affect the evaluation criterion.²

In this paper, we show that such a static description of people’s feelings of entitlement is incomplete. Our analysis is motivated by Kahneman et al. (1986), who argue that “when there is a history of transactions between firm and transactor, the most recent price, wage, or rent will be adopted for reference...” and that “terms of exchange that are initially seen as unfair may in time acquire the status of a reference transaction.”³ The hypothesis is related to the evidence on contrast effects in, e.g., judgments of the severity of crime (Pepitone and DiNubile, 1976), attractiveness ratings (Kenrick and Gutierrez, 1980), mate choice (Bhargava and Fisman, 2014) and financial markets (Hartzmark and Shue, 2015). We posit that such contrast effects operate in the domain of fairness as well. Consumers used to higher prices are more likely to perceive high prices as fair; employees used to high wages are more likely to perceive high wages as fair.

We also posit that such contrast effects stem from two broad types of experiences: *personal payoff experience*—the individually experienced payoffs resulting from a transaction— and *observational experience*—the outcomes observed from (possibly someone else’s) transaction.⁴ While observational experience and personal payoff experience are strongly correlated

¹For a review of the evidence on labor markets, see Fehr et al. (2009).

²See, for example, Rabin (1993), Fehr and Schmidt (1999), Bolton and Ockenfels (2000), Charness and Rabin (2002), Dufwenberg and Kirchsteiger (2004) and Falk and Fischbacher (2006). Models including reciprocity motives suggest a certain context dependence since the desire to treat someone kindly depends on how they acted. But the evaluation of an agent’s kindness still requires a static criterion.

³It has also been suggested that fairness judgments can also be affected by exposure to different kinds of bargaining environments (Binmore et al., 1991), or by investments into the production of the surplus that is to be divided between parties (Handgraaf et al. (2003) or Zwick and Mak (2012)).

⁴Similar differentiations have been made in the learning literature. For example, experience-weighted attraction learning (Camerer and Ho, 1999) assumes that more weight is given to directly experienced outcomes than to counter-factual outcomes when updating attractions of actions in a reinforcement learning type model.

in some types of transactions, they can substantially differ in other types of transactions such as posted offer markets with price discrimination, or in markets in which individuals sometimes act as sellers and sometimes as buyers. Consequently, a second goal of this paper is to quantify the importance of these two types of experience.

In our data—which involves a total of three experiments and 444 subjects—we find considerable path-dependence in fairness preferences. Subjects who have been exposed to unfavorable conditions towards themselves subsequently reveal less concern about being treated unfairly than subjects who have been exposed to favorable conditions towards themselves. We further find that observational experience accounts for more than half of this effect, though personal payoff experience plays an important role as well.

To identify a causal link from past experience to fairness perceptions, we conducted a two-phase experiment that exogenously varied experience in phase 1, and then used a simple game to elicit subject’s fairness preferences in phase 2. In the first phase, all subjects participated in one of two market games. In the *proposer competition* (PC) market (Roth et al., 1991), two *proposers* offer a monetary allocation to one responder, who can choose to accept either one or zero of those offers. In the *responder competition* (RC) market (Grosskopf, 2003), one proposer makes an offer to two responders, who simultaneously choose whether or not to accept the offer, with one responder randomly selected to transact in the case that both responders accept. Consistent with previous evidence, competitive pressures forced proposers in our experiment to give up most of their surplus in the PC market, while in the RC market, proposers kept most of their surplus.

In the second phase of the experiment, proposers and responders were matched one-on-one in a variant of the ultimatum game (Güth et al., 1982), and proposers again made offers to responders. Consistent with previous studies, responders were willing to reject an offer and forgo significant monetary gains to punish proposers making unfair offers. However, we find that responders’ experiences from the first part of the experiment had a significant impact on what offers they were willing to accept. In period 1 of the ultimatum game, the lowest acceptable offer of a responder who started in the PC market was on average 36% higher than the lowest acceptable offer of a responder who started in the RC market. That is, responders who started out in markets in which competition lead proposers to make very favorable offers to the responders had a much higher standard for what constitutes a fair and acceptable offer. We also find that this difference is persistent: over the course of 15 periods of repeated play, this difference dissipated by only about one-half of its period-one value.

To further quantify the deviation from static fairness models, we postulate a simple model of fairness preferences and estimate how its parameters depend on phase 1 outcomes.

Under different assumptions about which aspects of people’s fairness preferences are shaped by phase 1 experiences, our estimates imply that exposure to PC versus RC markets either changes perceptions of what’s fair by about 50%, or changes sensitivity to unfairness by about 50%.

To better understand how much of this effect is due to *personal payoff experience* vs. *observational experience*, we conducted two additional experiments. The *Role Switch experiment* was identical to the Baseline Experiment except that in phase 2, subjects who were previously proposers in phase 1 became responders in phase 2, and vice versa. This experiment is motivated by situations in which observational and personal payoff experience differ because an individual does not occupy a single role in a market. The *Full Information experiment* was identical to the Baseline Experiment except that in each round of phase 1, subjects received feedback not only about their own offer and payoff, but they were also informed about the average offer and acceptance rate in both markets. This experiment is motivated by the divergence between observed and experienced offers that can occur when the prices consumers pay for a particular good are different from the average price posted in a market, as in the case of publicly known price discrimination.

Combining the Baseline experiment with either the Role Switch or the Full Information experiment allows us to identify how much of the Baseline effect is driven by observational versus personal payoff experience. We find that both matter. Our results show that 50-75% of the path-dependence in fairness preferences is driven by observational experience, and 25-50% is driven by personal payoff experience.

Broadly, our results are consistent with the idea that contrast effects shape perceptions of fairness. In Section 4 we discuss possible theories that could generate our effects. We rule out standard theories of backwards-looking reference points for consumption bundles, “simple anchoring,” and (rational) expectations-based reference points. We discuss more likely explanations for our results: first, backwards-looking fairness reference points, and second, endogenous salience weights on fairness vs. payoffs (as would be microfounded by salience theory Bordalo et al. 2012, 2013, 2015).

Path dependence in people’s fairness preferences has immediate economic implications. It predicts inertia in how markets respond to changes in economic conditions. A sudden and major reduction in competition between sellers due to the exit of several competitors, for example, would not be followed by an equally sudden and drastic increase in posted prices – consumers used to low prices would not be willing to transact at significantly higher prices. Benjamin (2015) explores the theoretical implications of path-dependent fairness preferences in labor markets and shows how they can create downward wage stickiness, wage persistence

within a firm, and other empirical regularities documented in labor markets.⁵ Second, path-dependent fairness preferences generate new considerations for dynamic pricing strategies: posting a high price today would have the added benefit of increasing consumers’ willingness to pay a high price in the future. Third, our distinction between observational and personal payoff experience reveals new implications for managing price discrimination via information provision.

Generally, our paper is related to a nascent literature on preference formation and behavioral spillovers.⁶ Falk et al. (2006) show that minimum wage laws can cause spillover effects, raising wages even after the removal of the minimum wage law. Their results are consistent with the idea that minimum-wage policies are viewed as strong and salient “wage guidelines,” and provide evidence that people may infer social norms from government actions (e.g., Benabou et al. 2012). We differ from Falk et al. (2006) in that appropriateness of actions is not inferred from an authority; instead, we show that exogenous shocks to supply and demand shape subsequent perceptions of fairness through the channel of contrast effects.⁷ Others have shown that beliefs about opponents’ play can be influenced by observations of play in similar games, possibly through belief- or best-response bundling (Grimm and Mengel, 2012; Bednar et al., 2012; Cason et al., 2011).⁸ We, on the other hand, isolate the effect of past experience on subsequent judgments of fairness in a way that cannot be conflated with strategic considerations.⁹ More generally, and importantly, our paper goes beyond the

⁵See also Skott (2005) for an analysis of how fairness norms impact wage formation, and Kaur (2012) who formalizes the idea that workers may retaliate against a firm that offers them a wage below their reference wage.

⁶Recent evidence shows that social preferences can be affected by environmental determinants such as exposure to violence (see, e.g., Voors et al. (2012); Bauer et al. (2014); Cassar et al. (2012); Gilligan et al. (2014)), school intervention programs (Bettinger and Slonim, 2006) or class composition Rao (2013).

⁷Although Falk et al. (2006) designed their experiment to directly speak to minimum wage policies imposed by the government, the minimum wage policy in their experiment was of course imposed by the experimenter. The psychological mechanism, however, is essentially the same: subjects infer what is appropriate from the experimenter’s action. In fact, Falk et al. (2006) provide additional evidence that a simple suggestion from the experimenter about the appropriate wage has a significant impact on behavior. Additionally, see, e.g., Silverman et al. (2014) for evidence on subjects inferring appropriateness from experimenter’s actions.

⁸In the context of coordination games, Hargreaves-Heap and Varoufakis (2002) find that subjects who were previously disadvantaged in a Hawk-Dove game are more likely to coordinate on cooperative outcomes than subjects that were previously advantaged, and Lévy-Garboua et al. (2009) show that history can serve as focal point in multiple equilibria games. However, given the strategic nature of the interactions in these experiments, these data are also consistent with the possibility that subjects have standard preferences and best-respond to beliefs that are shaped by past experiences. An exception is the concurrent work by Peysakhovich and Rand (2015), who show that subjects who have previously experienced cooperative outcomes in the infinitely repeated prisoner’s dilemma not only act more cooperatively in strategic games, but also share more in the dictator game.

⁹For example, in Falk et al. (2006), the marginal revenue generated by a worker depends on the accept/reject decisions of other workers. In turn, the gains from trade and the distribution of surplus depend on other workers’ decisions. This implies that an individual worker’s behavior will not only be driven by his

existing literature on behavioral spillovers by distinguishing between the separate roles of observational and personal payoff experience, and by providing a direct and quantitative assessment of how the parameters of commonly used models of fairness preferences are shaped by past experiences.

Our work also contributes to a recent literature on the role of ex-ante agreements on the evaluation of ex-post outcomes in bilateral negotiations (Hart and Moore, 2008; Hart, 2009; Hart and Holmstrom, 2010; Fehr et al., 2011, 2015; Brandts et al., forthcoming), which argues that (potentially incomplete) contracts between two parties function as a reference point when evaluating the fairness of the final outcomes of the interaction between the two parties. Our notion of path-dependent fairness applies more broadly to environments in which parties do not have the opportunity to write a contract prior to choosing actions. And more importantly, we demonstrate that feelings of entitlement can endogenously be shaped by prior transactions *with other trading partners*, simply by interacting in a particular market environment.¹⁰

The rest of the paper is structured as follows. Section 2 presents our baseline experiment and results. Section 3 extends the experimental design to differentiate between observational experience and payoff experience. Section 4 discusses what theories can and can't explain our results. Section 5 concludes with a discussion of further applications and open questions.

2 Baseline Experimental Setting and the Path-dependent Fairness Hypothesis

2.1 Experimental design

All games in the experiment were based on the asymmetric ultimatum game, first introduced by Kagel et al. (1996)¹¹, and the market game first introduced by Roth et al. (1991). In each

own preferences, but also by his beliefs about the behavior of others.

¹⁰In fact, Hart and Moore (2008) discuss extensions of their model in which contrasts other than contractual terms affect parties' feelings of entitlement. Our work, therefore, paves the way toward more integrated models of reference-dependent fairness, that apply more broadly to not only interactions within bilateral trade agreements and organizations, but also within markets. An interesting insight from this strand of literature is that competitive mechanisms (Fehr et al., 2011) and bilateral negotiations Brandts et al. (forthcoming) are effective in shaping reference points, while exogenously assigned terms (Fehr et al., 2011) are not. This relates to evidence that the mere presence of a market can change people's motives to avoid moral transgressions or to engage in socially responsible behavior (Falk and Szech, 2013; Bartling et al., 2015a), and that competitive mechanisms affect perceptions of fairness violations (Bartling et al., 2015b). This points towards a potential interaction between the type of previous interaction and the outcome it generates in shaping path-dependent fairness preferences, which is an interesting avenue for future research.

¹¹This asymmetric ultimatum game is a variant of the original ultimatum game design first introduced by Güth et al. (1982).

of these games, 100 chips must be divided between proposers and responders, with proposers making offers, and responders choosing whether or not to accept the offers. These chips are then converted into monetary payoffs, with different conversion rates for the proposer and the responder. In our experiment, the monetary value of each chip was three times as high for a proposer as it was for a responder.¹² Our experimental design consists of three variants of the asymmetric ultimatum game: (i) Proposer Competition (PC), (ii) Responder Competition (RC) and (iii) no competition. Subjects participated in one of the two market games for the first 15 periods of our experiment, and then participated in the non-competitive ultimatum game in the next 15 periods. We describe the experimental games in more detail below.

2.1.1 Phase 1: Market Games

In the first phase of our experiment (first 15 periods), subjects participated in either a responder competition treatment or in a proposer competition treatment.

In the responder competition (RC) market game, one proposer is matched with two responders. The proposer first posts an offer of how to divide 100 chips between himself and a responder.¹³ Each responder then observes the offer and, without knowing the decision of the other responder, chooses whether or not to accept it. If both responders reject the offer, all three subjects receive zero chips. If one responder accepts the offer and one responder rejects the offer, the 100 chips are divided according to the proposed division between the proposer and the responder who accepted the offer. The responder who rejects the offer receives zero chips. If both responders accept the offer, it is randomly determined which responder actually receives the offer, and the non-selected responder receives zero chips.

In the proposer competition (PC) market game, two proposers are matched with one responder. Each proposer first posts an offer of how to divide 100 chips with the responder. The responder observes both offers and can accept one or none of the offers. If both offers are rejected, all three subjects receive zero chips. If an offer is accepted, the responder and the proposer who made receive chips according to the proposed split. The proposer whose offer was not accepted receives zero chips.

¹²We have chosen the asymmetric ultimatum game rather than the standard ultimatum game because existing evidence on responder behavior shows that the variance in minimum acceptable offers is considerably larger in the asymmetric ultimatum game than in the standard ultimatum game. Consequently, we considered the asymmetric ultimatum game to be better suited for treatment manipulations that seek to affect responder behavior.

¹³In all games, offers had to be multiples of 5 chips.

2.1.2 Phase 2: Ultimatum Game

In the next phase of our experiment (next 15 periods), all subjects participated in a standard version of the asymmetric ultimatum game for 15 periods. In this version, one proposer is matched with one responder. First, the proposer makes an offer to the responder. Second, the responder can accept or reject the offer. We did not elicit responders' decisions in phase 2 in the same way that we elicited them in phase 1. Before responders are informed about the actual offer, but after the offer is made, responders state a minimum acceptable offer (MAO) amount; that is, each responder states a number x such that the proposer's offer is accepted if and only if he offers at least x chips to the responder. This minimum amount is binding and directly enforced by the computer. As before, the proposed division of chips is implemented if and only if the proposer's offer is accepted, while both subjects received zero chips if the proposed offer is rejected.¹⁴ Importantly, proposers are never informed ex-post about the responder's MAO, but only about whether the offer is accepted or rejected.

2.1.3 Procedures

At the beginning of each session, each subject was assigned to the role of proposer or responder, and this role was fixed throughout the experiment. Just before the first period, one third of the proposers and two thirds of the responders were randomly assigned to the proposer competition treatment. The remaining two thirds of the proposers and one third of the responders were assigned to the responder competition treatment. Subjects stayed in their respective treatment groups throughout all of phase 1 of the experiment. All subjects received written instructions for their respective treatment, and were asked to answer several understanding checks before proceeding with the experiment. After all subjects completed the instructions and the understanding checks, they were asked to proceed to the first phase of the experiment. Proposers and responders were randomly re-matched within their treatment group after every period. The subjects were told that there would be a second phase to the experiment, but were told nothing else about it other than that their choices in phase 1 would have no effect on their potential payoffs in phase 2.

Once the first phase of the experiment was finished, subjects received on-screen instructions for the ultimatum game without competition, and were again asked to work through

¹⁴Our use of the strategy method in phase 2 but not in phase 1 implies a difference in the responders' choice sets between the two phases. In phase 1, responders are given choices $A_1 = \{accept, reject\}$, while in phase 2, they are given choices $A_2 = \{0, 5, \dots, 100\}$. We did not use this strategy in phase 1, because we didn't want to exogenously impose rules about which offer must be chosen under proposer competition. Also, note that eliciting MAO's is technically not fully equivalent to the strategy method, since a responder's full strategy might be to accept an offer of x but reject an offer $y > x$. But as long as responders' acceptance preferences are monotonic, there is no loss of information in eliciting MAOs.

several understanding checks. They were then divided into three different matching groups. Each matching group contained one third of the proposers and one third of the responders within a session. The first matching group consisted of proposers and responders who had previously been in the proposer competition treatment (PC Matching Group). The second matching group consisted of proposers and responders who had previously been in the responder competition treatment (RC Matching Group). Finally, the third matching group consisted of the remaining third of proposers who had previously been in the proposer competition treatment and the remaining third of responders who had previously been in the responder competition treatment (Mixed Matching Group).¹⁵

Table 1: Overview of Matching Groups

	Proposer Phase 1 Experience	Responder Phase 1 Experience
PC Matching Group	PC Proposers	PC Responders
RC Matching Group	RC Proposers	RC Responders
Mixed Matching Group	PC Proposers	RC Responders

As a naming convention, we will refer to responders and proposers who have previously participated in the proposer competition market as “PC Responders” and “PC Proposers”, and to those who have participated in the responder competition market as “RC Responders” and “RC Proposers”. The composition of the matching groups is summarized in table 1. Subjects stayed within their respective matching groups throughout all 15 periods, but in every period responders and proposers were randomly rematched within their matching group, to rule out strategic incentives in the choice of the MAO.¹⁶ The matching groups allow us to cleanly investigate the effect of responder experience on bargaining behavior, holding proposer experience constant.

To avoid wealth effects potentially confounding or interfering with our treatment manipulation, either phase 1 or phase 2 was selected for payment at the end of the experiment.¹⁷ Within the chosen phase, 4 periods were selected at random.¹⁸ The points earned in the selected periods were then converted into Swiss Francs, with the exchange rate of points to

¹⁵Subjects were fully informed about all aspects of the games they directly participated in, but they were not informed about the respective other market game. In phase 2, subjects were therefore not informed about the market history of their matches.

¹⁶We therefore implemented a stranger and not a perfect stranger matching protocol. However, evidence suggests that subjects treat these protocols similarly. In their meta-analysis of ultimatum game behavior, Cooper and Dutcher (2011) find no evidence of different behavior between stranger and perfect stranger protocols.

¹⁷In an expected utility framework, the independence axiom implies that our payoff structure guarantees that phase 1 history should be irrelevant to people’s preferences over phase 2 outcomes.

¹⁸We selected 4 periods rather than 1 to reduce the variance in subject payments in case phase 1 of the experiment was selected for payment (which otherwise would have been very large).

Swiss Francs set at 10:1.

In total, we ran 5 sessions of the Baseline Experiment, totaling to 150 subjects.¹⁹ Experiments were computerized using the software z-tree (Fischbacher, 2007) and conducted at the experimental laboratory of the University of Zurich. Our subject pool consisted primarily of students at the University of Zurich and the Federal Institute of Technology in Zurich.²⁰ On average, an experimental session lasted 75 minutes with an average payment of CHF 38.5 (\$42.00), including a show-up fee of CHF 10.²¹

2.2 Conceptual Framework for the Path-dependent Fairness Hypothesis

We formalize our hypotheses with an intentionally simple extension of the well-known social preference models introduced by Fehr and Schmidt (1999; henceforth FS), Bolton and Ockenfels (2000, henceforth BO), and Charness and Rabin (2002, henceforth CR). Like, FS, BO, CR, we capture several key properties of fairness preferences using a maximally tractable model.

We consider an N -player game in which we let π_i denote each player’s final monetary payoff. We let player i ’s utility be given by

$$U_i = \pi_i - \beta(h)\max(r(h)\Pi - \pi_i, 0) - \alpha(h)\max(\pi_i - r(h)\Pi, 0),$$

where $\Pi = \sum_j \pi_j$ is the total surplus, and h is the experienced and/or observed history of offers and payoffs. Here, r is the share of the total surplus a player feels entitled to, or what BO call the “perceived social reference point”. Concretely, our model is a piecewise-linear version of BO’s ERC model. The parameters $\alpha \in [0, 1]$ and $\beta > \alpha$ capture, respectively, the disutility associated with player i feeling that he got more or less than what he feels he is entitled to.²²

Both FS and BO focus on equity theories in which perceptions of fairness are based on a fixed, exogenously given equity norm that is not shaped by past experience. In the context

¹⁹Because differences in past experience are a crucial variable in our design, we only invited subjects who have not previously participated in ultimatum game experiments.

²⁰Subjects were drawn from a database of volunteers using ORSEE Greiner (2004).

²¹In all sessions, we also elicited beliefs of proposers and responders about average offers and average acceptable offers. Moreover, at the end of the experiment, all subjects participated in the cognitive reflection test (Frederick, 2005). Subjects received an additional CHF 5 for these tests. We report results with respect to these measures in an earlier working paper version of this paper Herz and Taubinsky (2013). Because they are not essential for any of our results, we do not further discuss these measures in this paper.

²²Following CR and others, we make the assumption $\alpha \leq 1$ to capture the idea that a player won’t ever burn Δ of his money just so he doesn’t get more than his fair share.

of our specific formal model, this would translate into $r(h) = 1/N$ for all players in an N -player game, irrespective of past experiences. We, however, do not set $r(h)$ equal to $1/N$ but instead allow for the possibility that it may be shaped by past experience h .²³

Similarly, FS and BO assume fixed and exogenous preference parameters α and β . Alternatively, the sensitivity to (un)fairness may be history-dependent. One mechanism for this is shifts in salience based on previous experiences, as in, e.g., Bordalo et al. (2015). We provide this microfoundation in Appendix C.

What are the testable implications of path-dependent fairness preferences? For phase 1 of our experiment, the possibility that r , α and β are potentially shaped by past experience does not generate sharp testable implications. In Appendix A, we generalize the FS and BO theoretical results about fairness and market competition, and show that in our more general framework, proposer competition still drives proposer surplus to zero, while responder competition drives responder surplus to zero.

In phase 2, behavior is much more sensitive to the social referent r and the sensitivity to negative inequality β . In the context of our experimental payoffs, simple algebra shows that the smallest offer a responder is willing to accept is given by

$$MAO(r) = \frac{300\beta r}{2\beta r + \beta + 1}, \quad (1)$$

which is a strictly increasing function of r and β . Thus if experience affects either of these parameters, then it should have a direct effect on the minimally acceptable offers (MAO's) of responders.

Our basic hypothesis is that responders who are used to receiving low offers in phase 1 will have lower minimal acceptable offers than responders who are used to receiving relatively high offers from proposers, either because of a lower reference point or because fairness concerns are less salient to them.

Our analysis of phase 2 behavior will focus on responders because their behavior is solely a function of the preference parameters β and r , rather than strategic considerations about other players' behavior. Proposers' behavior, by contrast, is shaped by their beliefs about responder behavior, in addition to their social preferences.²⁴ We thus focus most of the analysis on responders' MAO's, but return to exploring proposer behavior in Section 3.5.

²³In principle, $r(h)$ could be a function of more than just past experience. As noted in footnote 3, entitlements have also been shown to be affected by ex-ante investments into the production of the surplus that is to be divided. Our framework could be amended to also capture such influences. For simplicity and tractability, however, we will solely focus on past experience as a determinant of fairness preferences.

²⁴Letting $Q(a)$ denote a proposer's belief that his offer a will be accepted, the proposer chooses a to maximize

$$3(100 - a)Q(a) - \beta \max[r(300 - 2a) - 3(100 - a), 0] - \alpha \max[r(300 - 2a) - a, 0] \quad (2)$$

2.3 Behavior in the Phase 1 Market Treatment

Our phase 1 treatment variation was successful in inducing large, exogenous differences in phase 1 experience. As expected, competition had a strong effect on offers in the first phase of our experiment. Averaged over all 15 periods, proposers offered 78 chips to responders in the PC market, whereas they offered only 31 chips to responders in the RC market. The development of offers over the course of the 15 periods in both treatments is shown in the left panel of figure 1. The difference between offers in the two treatments is roughly 23 chips in period 1, and increases over time until it reaches an average of 50 chips from period 7 onwards. The average difference in offers between the markets is 46.7 chips, and this difference is highly statistically significant in a regression of offers on a PC market dummy, with standard errors clustered at the level of phase 1 market / session pairs.

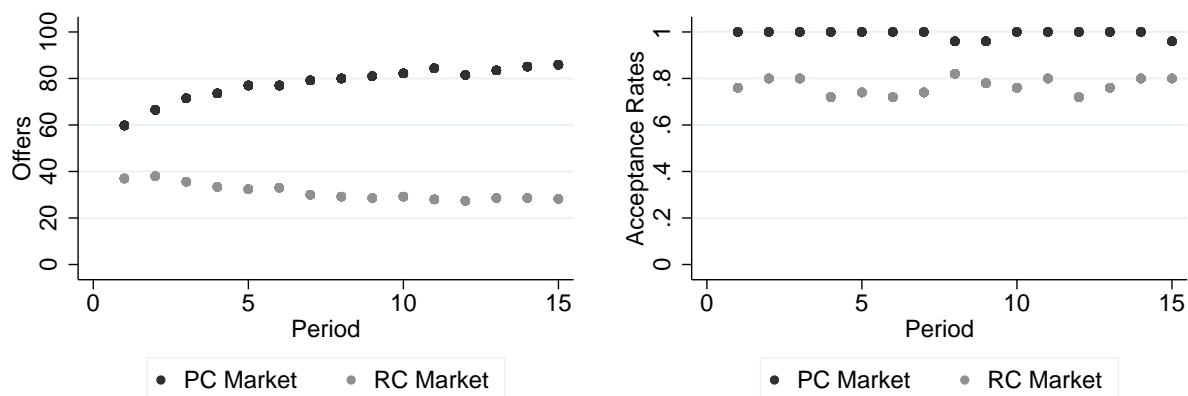


Figure 1: Left Panel: Average offers over time under responder competition (RC Market) and under proposer competition (PC Market) in phase 1 of the Baseline experiment. Right Panel: Acceptance rates of responders over time under responder competition and under proposer competition in phase 1 of the Baseline experiment.

The right panel of figure 1 shows that the probability that an offer was accepted. In the PC market, responders accept one of the two offers in 99.2 percent of the time. In the RC market, responders accept the offers 76.8 percent of the time, and the probability that *at least one* of the responders accepts an offer is 92.5 percent. Thus in both markets, a successful transaction occurs over 90 percent of the time. Our stark results on the effects of competitive forces are consistent with Roth et al. (1991), Grosskopf (2003) and Fischbacher

Equation 2 shows that interpreting the impact of phase 1 experience on proposers' offers can be problematic for two reasons: First, it is unclear whether phase 1 experience affects $Q(a)$ or one of the preference parameters. Second, all models of fairness assume that people are more concerned about being "behind" than about being "ahead"; i.e., $\alpha < \beta$. But if α is small relative to β , then changes in r will have a smaller impact on proposers' offers than on responders' MAO's, and changes in β are not detectable at all in proposer behavior.

et al. (2009). Our two market treatments thus generate substantial exogenous variation in phase 1 experience for testing the path-dependent fairness hypothesis in phase 2.

2.4 Behavior in Phase 2: The Effect of Phase 1 Experience on Responder Behavior in Phase 2

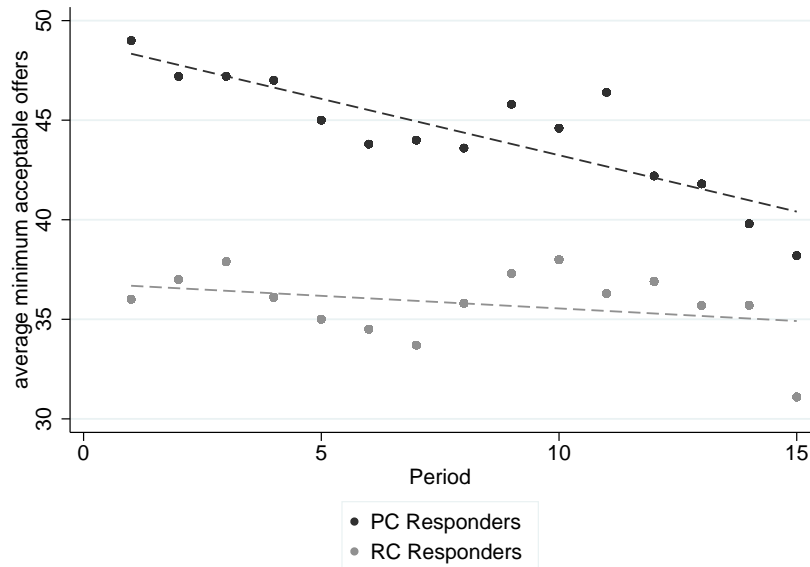


Figure 2: Minimum Acceptable Offers of responders. “PC Responders” denotes responders who have participated in the PC market in phase 1. “RC Responders” denotes responders who have participated in the RC market in phase 1. The figure shows average minimum acceptable offers for PC responders and RC responders over the course of the second part of the experiment. The dashed lines show the linear time trends.

Figure 2 plots responders’ minimal acceptable offers. In every period of phase 2, average minimal acceptable offers are larger for PC Responders, and the difference is particularly pronounced in early periods.

To quantify the effect of phase 1 experience on responder MAO’s, we estimate OLS regressions of individual MAO’s on phase 1 market dummies. Table 2 shows results of such regressions. Standard errors are clustered at the level of the phase 1 market matching group. Column (1) shows the average treatment effect in period 1, whereas column (2) shows the average treatment effect over all 15 periods. Column (1) shows that responders that have previously been in the proposer competition market (indicated by the dummy “PC Responder”) have minimum acceptable offers that are 13 chips higher than responders who have previously been in the responder competition market, which translates to PC

Table 2: The Impact of Responder Experience on minimal acceptable offers (MAO) in the Baseline Experiment

	(1)	(2)
	MAO	MAO
PC Responder	13.00** (4.21)	10.19* (4.92)
PC Proposer		-3.23 (7.49)
Constant	36.00*** (2.90)	37.41*** (4.68)
Adj. R^2	0.06	0.04
Observations	75	1125

The regression in column (1) includes observations from period 1 only. The regression in column (2) includes observations from all periods. PC Responder is a dummy variable indicating whether a responder participated in the PC market in phase 1. PC proposer is a dummy variable indicating whether the matched proposers participated in the PC market in phase 1. PC Proposer is only controlled for in column (2), since responders had no interaction with their matches in phase 2 prior to entering their period 1 MAO. Controlling for PC Proposer in column (1) leads to insignificance of the PC Proposer dummy and leaves significance of the PC Responder dummy unchanged (regressions not reported here). Robust standard errors are clustered by phase 1 market matching groups (2 clusters per session, 10 clusters in total). Significance levels: *** = 1%, ** = 5% and * = 10%.

Responders stating minimum acceptable offers that are 36 percent higher than the acceptable offers of RC Responders. This difference is significant at the 5% level. Column (2) shows that the effect of phase 1 experience remains significant at the 10% level even when all 15 periods are considered.

Figure 3 plots regression coefficients corresponding to the difference in MAO’s between PC and RC responders in each period.²⁵ The figure shows that the effect of phase 1 experience decays only slightly to about 10 chips over the course of the 15 periods.²⁶ Intuitively, differences in MAO’s in period 1 of phase 2 capture the direct effect of the exogenous variation in phase 1 experience. In periods $t > 1$ in phase 2, however, both RC and PC responders begin to play the same game, and thus their experiences begin to become more similar. Consequently, their history-dependent fairness preferences are expected to converge.²⁷

²⁵15 regressions identical to the regression in column (2) of table 2 were conducted, one for each period. The figure shows the coefficient on the PC Responder dummy.

²⁶We also ran a regression interacting the PC Responder dummy with period, to estimate a linear time trend. We find that the coefficient on PC Responder is decreasing by 0.44 percentage points per period, and this negative time trend is significant at the 10% level.

²⁷We demonstrate this point formally for the case of path-dependent fairness reference points in Appendix D.

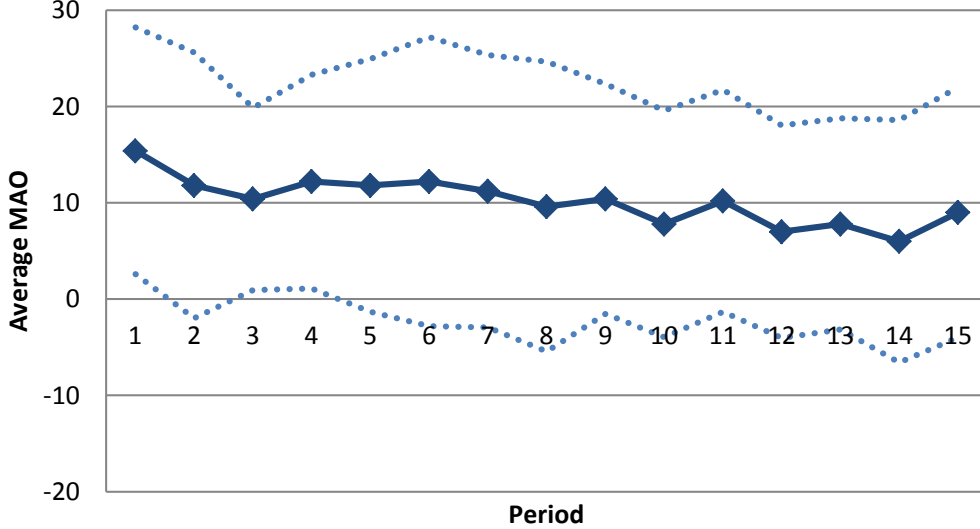


Figure 3: This figure shows the development of the coefficient on PC responder over time, individually estimated for every period using the same regression specification as in column (1) of table 2. The regressions also include PC proposer dummies, to control for any potential impact of proposer experience. However, as already evident in column (2) of table 2, these dummies are insignificant. Performing the same regressions without these dummies does not significantly alter the results. The dotted lines indicate 95% confidence intervals. Robust standard errors are clustered by the phase 1 market treatment group (2 clusters per session, 10 clusters in total).

2.5 Impacts of Phase 1 Experience on Parameters of the Fairness Model

The MAO analysis provides clear evidence of path-dependence of social preferences, but it does not provide clear guidance about the magnitude of path-dependence in terms of the parameters of fairness models. To quantify how much the parameters in classical fairness models such as those of FS, BO, and CR can be shaped by past experience, we now estimate how phase 1 experience shapes either the phase 2 fairness reference point r or the sensitivity to negative inequality β in equation (1). As with our reduced-form results, our goal with the structural estimation strategy is to utilize only exogenous variation created by our two phase 1 treatments, rather than to utilize all (potentially endogenous) variation in MAO's in phases 1 and 2. The equations that utilize only treatment-level differences to identify the parameters are:

$$MAO_{j,it} = \frac{300\beta_j r_j}{2\beta_j r_j + \beta_j + 1} + \epsilon_{it} \quad (3)$$

Table 3: Structural Estimates of fairness preference parameters in the Baseline Experiment

	Estimates for r		Estimates for β	
	(1)	(2)	(3)	(4)
ρ	0.11*** (0.03)	0.07** (0.03)		
β	0.67*** (0.07)	0.58*** (0.06)		
β_{RC}			0.46*** (0.07)	0.46*** (0.04)
β_{PC}			0.94*** (0.16)	0.72*** (0.14)
Observations	75	1125	75	1125

Columns (1) and (2) contain estimates for ρ and β , where $r_{PC} = 0.5 + \rho$ and $r_{RC} = 0.5 - \rho$. Column (1) uses period 1 data only. Column (2) uses data from all periods. Column (3) contains estimates of β_{RC} and β_{PC} using period 1 data only, whereas column (4) uses data from all periods. Robust standard errors are clustered by phase 1 market matching groups (2 clusters per session, 10 clusters in total). Significance levels: *** = 1%, ** = 5% and * = 10%.

where $j = \{RC, PC\}$, the subscript RC is used for parameters of RC responders, the subscript PC is used for parameters of PC responders, $MAO_{RC,it}$ and $MAO_{PC,it}$ are the MAOs of a responder i in period t in RC and PC markets, respectively, and ϵ_{it} is a mean zero person-period specific error term. Reformulating and setting $E[\epsilon_{it}] = 0$, our exogenous variation in phase 1 market experience then gives us two moment conditions, containing a total of four parameters ($\beta_{RC}, r_{RC}, \beta_{PC}, r_{PC}$) that cannot be identified simultaneously. Intuitively, our data cannot distinguish whether MAOs are higher because the reference point r is higher or because the sensitivity parameter β is higher. We thus consider two different simplifications of the model.

First, we assume that $\beta_{RC} = \beta_{PC} \equiv \beta$, and estimate how the reference point r would have to change to accommodate the observed behavior. This leaves us with three parameters, which we reduce to two by setting $r_{RC} = 1/2 - \rho$ and $r_{PC} = 1/2 + \rho$.

Second we set $r_{PC} = r_{RC} = 1/2$, and estimate how the sensitivity parameters β_{RC} and β_{PC} must differ to accommodate the observed behavior. We perform our estimation using the method of moments. Given the two-parameter vector $\xi = (\beta, \rho)$ or $\xi = (\beta_{RC}, \beta_{PC})$, let $m(\xi)$ denote the theoretical vector of moments corresponding to the two equations above. Because our model is *exactly identified*, the estimation procedure here is simple: for observed moments \hat{m} , the method of moments estimator chooses the parameter vector $\hat{\xi}$ for which $m(\hat{\xi}) = \hat{m}$. As with the reduced-form regressions, we compute robust standard errors clustered at the phase 1 market level for each session.

Table 3 presents our results, computed using the Gauss-Newton algorithm. Column (1)

estimates ρ and β using data from period 1 only, while column (2) estimates ρ and β using data from all 15 periods. We note that the estimates in column (2) do not capture the full dynamics of preferences, which may evolve period by period over the course of phase 2. However, because our exogenous variation comes from Phase 1 only, we do not have sufficient variation to identify the full dynamics over the course of Phase 2. Rather, the question we ask is what fixed pair (r_{RC}, r_{PC}) best rationalizes the aggregate behavior in Phase 2. In section 3, however, we will specify a functional form for the evolution of preferences over time, and estimate the parametrized dynamics. Columns (3) and (4) estimate β_{RC} and β_{PC} , again using either only period 1 data (column (3)) or the data from all 15 periods (column(4)).

Column (1) shows an estimate of $\rho = 0.1$, suggesting that exogenous variation in experience as extreme as the difference between the two markets can change the entitlement r by about 30%-50%: RC market experience generates $r \approx 0.4$, while PC market experience generates $r \approx 0.6$. This difference in r is significant at $p < 0.01$. When using data from all periods in column (2), the difference is smaller, but still significant at $p < 0.05$. That the differences become attenuated over time is not surprising, and arises naturally from a dynamic extension of our model, presented in Appendix D. These structural estimates show that the deviation from FS, BO, and CR type models—in which there is an exogenous equity norm $r = 1/2$ in a two player game like phase 2 of our experiment—is not only statistically significant, but also economically significant and large in magnitude.

If we fix the fairness reference point instead, and assume that our treatment affects responders sensitivity to negative inequality β , we find similarly stark results. Column (3), which uses only period 1 data, shows that $\beta_{PC} = 0.94$, which is more than twice as large as $\beta_{RC} = 0.46$. The difference between these estimates is significant at $p < 0.01$. When we use data from all 15 periods, the estimate for β_{PC} decreases modestly to 0.72, while the estimate for β_{RC} remains largely unchanged at 0.46. The difference in these estimates is still significant at $p = 0.059$.

3 Unpacking the Channels of Experience Effects

3.1 Conceptual Framework

Our baseline experiment demonstrates that responders' preferences are influenced by phase 1 experiences. However, our experimental design in phase 1 does not allow us to differentiate between the possible influences of two broad types of experiences generated by our market games: *observational experience* and *personal payoff experience*.

We define observational experience as the average offer observed in the phase 1 market.

This type of observation is independent of an individual’s role in the market, the specific offer at which an individual transacted, and the extent to which an individual even participated in the market at all.²⁸

Additionally, preferences may also be shaped by one’s own, specific personal payoff experience. Regardless of what is the average price of some good, what a person considers a reasonable price to buy at may depend on the specific share of the surplus that he had been receiving in previous transactions. A person used to buying at low prices may feel averse to buying at high prices because he is not used to giving up most of the transaction surplus. We call this type of experience *personal payoff experience*, and it is defined as the average share of the group’s total payoff in each round of the market or ultimatum games.

In our Baseline Experiment, these two components of experience are almost perfectly correlated. Proposer competition generates high offers, which generates both high observational and high personal payoff experience for responders. To separate the potentially differential impact of observational and personal payoff experience, we therefore ran two additional experiments, the Role Switch experiment and the Full Information experiment. Both experiments are motivated by two economically meaningful ways in which observational experience and personal payoff experience differ. First, they may differ substantially for individuals who don’t occupy a single role in a market. Consider previous employees who worked for low wages and transition to the role of employer and compare them to employers in a low wage market who transition to the role of an employee. They will observe the same market outcomes—they both participated in a market with low wages—but their personal payoff experiences will likely differ. Our Role Switch experiment is motivated by this kind of role reversal.

Our Full Information experiment is motivated by situations in which observational and personal payoff experience can differ substantially when the price a consumer pays for a particular good is different from the average price posted on the market. Suppose, for example, that a consumer faces price discrimination, but can fully observe the menu of prices that a monopolist is offering on the market. This price-discriminating monopolist sells a widget for \$10 to one group of consumers and for \$20 to another group of consumers, and this price-discrimination is known to everyone in the market. Here, both the \$10 and \$20 consumers will observe the same set of prices but, by definition, they will have different personal payoff experiences.

²⁸In principle, observed acceptance or rejection could matter as well. However, acceptance rates in both market games are very high, and we do not have exogenous variation in acceptance rates conditional on offers. Hence, we focus on observed offers in our analysis. The impact of acceptance on fairness reference points may be an interesting avenue for future research.

3.2 Design of the Role Switch and the Full Information Experiments

3.2.1 The Role Switch Experiment

Phase 1 of the Role Switch experiment was identical to phase 1 of the Baseline experiment. Once phase 1 was finished, subjects were presented with a new set of instructions for phase 2 of the experiment. However, in the Role Switch experiment all subjects that were assigned to the responder role in phase 1 were re-assigned to the proposer role in phase 2, and all subjects that were assigned to the proposer role in phase 1 were re-assigned to the responder role in phase 2. The role switch reverses the correlation between personal payoff experience and observational experience relative to the Baseline experiment, since proposers in the RC market observe low offers, but receive high payment shares, and vice versa. Other than the role reassignment, phase 2 was equivalent to phase 2 of the Baseline experiment.

In total, 4 sessions of the Role Switch experiment were conducted. 30 subjects participated in each session, leading to a total of 120 subjects who participated in the Role Switch experiment. Sessions lasted approximately 1-1.25h and subjects on average earned 36.4 CHF (approx. 40 USD) including a 10 CHF show up fee.²⁹

3.2.2 The Full Information Experiment

The Full Information experiment differed from the Baseline experiment in the feedback given to subjects during phase 1. After every period of phase 1, all subjects were informed about the *average offer* as well as the *average acceptance rate* in both the PC market and the RC market.³⁰ Consequently, in the Full Information experiment, observational experience is held constant for all subjects, independent of the market they have been assigned to in phase 1. Phase 2 of the Full Information experiment was again similar to phase 2 of the Baseline experiment.

²⁹In the Role Switch experiment, subjects' offer screen featured lists that included all possible offers respectively minimal acceptable offers. This is a difference to the Baseline experiment, in which proposers simply entered numbers. We did this to actually make our phase 2 responders' decision format *more* comparable to the format in the Baseline experiment. In the baseline experiment, responders first made binary decisions in phase 1, and then selected an MAO in phase 2. This design choice minimized the possibility of mindless anchoring, in the sense that subjects might simply continue entering the same number over and over again, irrespective of what phase of the experiment they're in. In the Role Switch experiment, we similarly wanted the phase 2 responders (who are phase 1 proposers) to make binary choices in phase 1, so as to minimize the possibility of mindless anchoring on phase 1 choices. This slight change of format did not alter the phase 1 behavior, as shown in Appendix I.

³⁰In contrast to the other two experiments, subjects were therefore fully informed about the two simultaneously conducted markets in phase 1 of the experiment. During the experiment and in the instructions, the two different types of markets were not referred to as "proposer competition market" and "responder competition market", but as "market of type X" and "market of type Y", respectively.

In total, 6 sessions of the Full Information experiment were conducted. 24 or 30 subjects participated in each session, leading to a total of 174 subjects who participated in the Full Information experiment. Sessions lasted approximately 1-1.25 h and subjects on average earned 36.75 CHF (approx. 41 USD), including a 10 CHF show up fee.

3.2.3 The Six Experimental Conditions

Table 4: Overview of treatment variation for phase 2 responders across all three experiments

	Observational experience		
	High	High&Low	Low
High personal payoff experience	PC Responders (Baseline)	PC Responders (Full Info)	RC Proposers (Role Switch)
Low personal payoff experience	PC Proposers (Role Switch)	RC Responders (Full Info)	RC Responders (Baseline)

Table 4 summarizes how our three experiments allow us to separately identify the effects of observational experience and personal payoff experience. Roughly, our three experiments generate six different cells: (high vs. low personal payoff experience) \times (high vs. high & low vs. low observational experience). PC responders in the Baseline experiment and in the Full Information experiment, as well as RC proposers in the Role Switch experiment are categorized into “high personal payoff experience”, whereas RC responders in the Baseline experiment and the Full Information experiment, as well as PC proposers in the Role Switch experiment are categorized into “low personal payoff experience”. PC responders in the Baseline experiment as well as PC proposers in the Role Switch experiment are categorized into “high observational experience”, and RC responders in the Baseline experiment as well as RC proposers in the Role Switch experiment are categorized into “low observational experience”. All subjects from the Full Information experiment are categorized as having observed both high & low offers.

Table 5 summarizes the average observational and personal payoff experience in phase 1 for each treatment group in each experiment. The table shows that our exogenous treatment variation indeed had the expected effects.

3.3 Impact of Payoff and Observational Experience on MAOs

To formalize the notion of personal payoff experience, let $\mu_i^t = \pi_i^t/\Pi^t$ denote the share of period t surplus that player i receives, and let $\bar{\mu}^\tau = \sum_{t=1}^{\tau-1} \mu^t/(\tau - 1)$ denote the average of these experiences. In the case that $\Pi^t = 0$, we set $\mu_i^t = 0$, to reflect our intuition that

Table 5: Summary statistics of observational experience and personal payoff experience by treatment and experiment

	Personal payoff experience	Observational experience
Baseline PC Responders	0.63	0.57
Baseline RC Responders	0.07	0.14
Role Switch PC Proposers	0.19	0.58
Role Switch RC Proposers	0.70	0.22
Full Information PC Responders	0.66	0.43
Full Information RC Responders	0.07	0.43

receiving a zero payoff should lower one’s feelings of entitlement.³¹

Second, we formalize observational experience. Given offers $a_1^t, a_2^t, \dots, a_n^t$ observed by some player j in period t , let $\nu_j^t = \frac{1}{n} \sum_i \frac{a_i^t}{\pi_P(a_i^t) + a_i^t}$ denote the average normalized offer observed by player j . Note that the quantity $\frac{a_i^t}{\pi_P(a_i^t) + a_i^t}$ has a simple interpretation: it is the share of the surplus that has been offered. Let $\bar{\nu}^\tau = \sum_{t=1}^{\tau-1} \nu^t / (\tau - 1)$ denote the average of these observations.

In this section, we examine the impact that $\bar{\mu}^\tau$ and $\bar{\nu}^\tau$ have on phase 2 behavior. Although $\bar{\mu}^\tau$ and $\bar{\nu}^\tau$ are very simple summary statistics of experience, in appendix F we also show that our empirical results are robust to considering weighted averages of past experiences and observations: $\bar{\mu}^\tau = \sum \delta^t \mu^t / \sum \delta^t$ and $\bar{\nu}^\tau = \sum \delta^t \nu^t / \sum \delta^t$.

We begin by summarizing the behavior conditional on the treatment.³² Figures 4 and 5 show the average MAOs for the different combinations of observational and personal payoff experience, as described in table 4. Figure 4 summarizes the data in a way that makes transparent the impact of observational experience. The top panel shows the effect of observational experience, pooling data across low and high personal payoff experience. The two bottom panels show the effect of observational experience, conditional on low (panel (b)) and high (panel (c)) personal payoff experience. Figure 5, on the other hand, summarizes the data in a way that makes transparent the impact of personal payoff experience. The top panel shows the average effect of personal payoff experience, pooling across all experimental conditions. The three bottom panels show the effect of personal payoff experience, condi-

³¹An alternative intuition is that in an N -player group, $\mu_i^t = 1/N$ when $\Pi^t = 0$, to reflect the possibility that when everyone gets the same payoff (even when it’s zero), the player feels like it was such an equitable outcome that his subsequent feelings of entitlement move towards him getting an even share of the surplus. In Appendix F, we show that our results are nearly identical under this alternative specification.

³²In Appendix I, we show that the phase 1 outcomes in these two additional experiments are very similar to the phase 1 outcomes in the Baseline experiment. Market forces work as expected and drive up offers in the PC markets, whereas they drive down offers in the RC markets, again leading to substantial exogenous variation in experiences. In the rest of our analysis, we thus focus on phase 2 only.

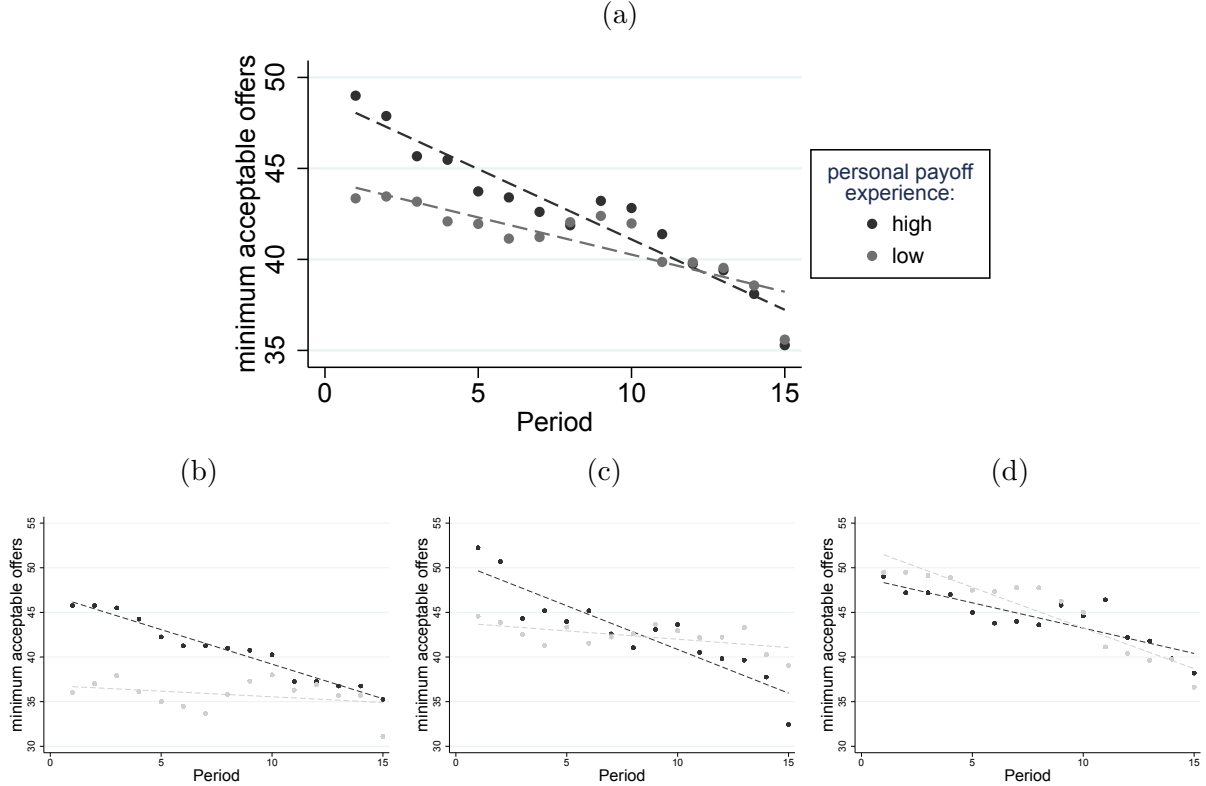


Figure 4: Minimal Acceptable offers by personal payoff experience. The top panel (a) pools data from all treatments, independent of observational experience. The bottom panels condition on either low personal payoff experience (b), high and low observational experience (c) or high personal payoff experience (d). See table 4 for the respective treatment descriptions.

tional on low (panel (b)), high and low (panel (c)), and high observational experience (panel (d)).

To statistically analyze the effect of personal payoff experience and observational experience on responder MAO's separately, we first analyze the Baseline and the Role Switch experiments in isolation, and then analyze the Baseline and Full Information experiments in isolation. We estimate instrumental variables regressions using our six treatment conditions (the random assignment to the RC and the PC market in each of our 3 experiments) as instruments for responder personal payoff experience and observational experience. We use instrumental variable regressions instead of OLS regressions because there is simultaneity bias at the matching group level: Responders' and proposers' preferences shape the outcomes in phase 1 in their respective matching groups, and those outcomes will be related to Phase 2 outcomes not just through the causal experience channel, but also simply because of within-subject—and thus within-matching-group—correlation in behavior.³³

³³To put in another way: for the same reason that we can not test our path-dependence hypothesis by

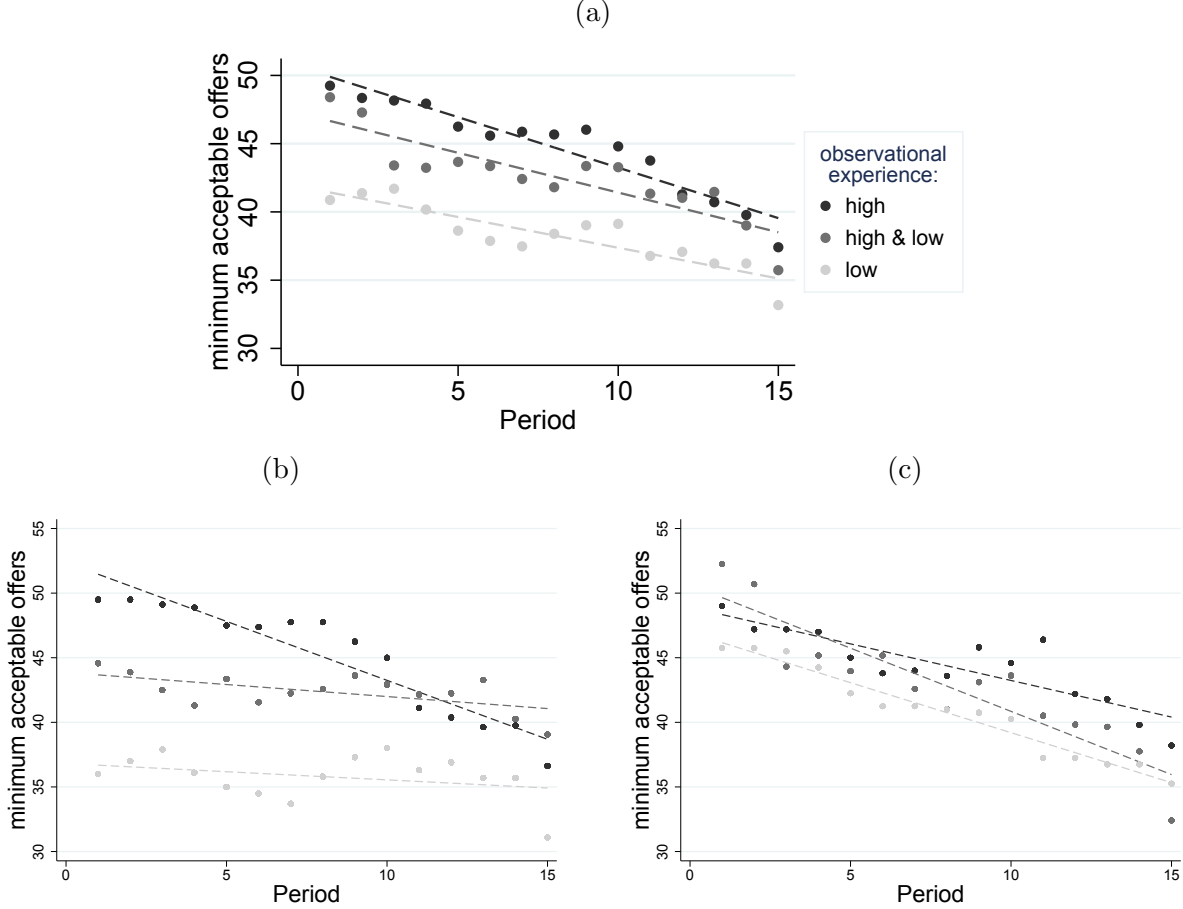


Figure 5: Minimal Acceptable offers by observational experience. The top panel (a) pools data from all treatments, independent of personal payoff experience. The bottom panels condition on either low personal payoff experience (b) or high personal payoff experience (c). See table 4 for the respective treatment descriptions.

Because our standard errors are not homoscedastic, we use the more efficient iterative Generalized Method of Moments (GMM) estimator (Hall, 2005) instead of the two-stage least squares (2SLS) estimator. Our 6 moment conditions are $E[(MAO - \beta_0 - \beta_1\bar{\mu} - \beta_2\bar{\nu})T_j] = 0$, where T_j is a dummy variable for one of the six phase 1 treatment conditions and $\bar{\mu} = \sum_{t=1}^{15} \mu^t / 15$ and $\bar{\nu} = \sum_{t=1}^{15} \nu^t / 15$ are the average of phase 1 observational experience and personal payoff experience. Our estimates here do not account for the fact that our framework predicts that preferences should be changing throughout phase 2 as well. Rather, the question we ask is how do phase 1 observational and payoff experiences affect behavior in phase 2, on

simply having subjects play 15 rounds of an ultimatum game and then regressing their MAO's on past experience, we can't simply regress MAOs on past experience even when we have an additional source of true exogenous variation. The instrumental variables regression allow us to focus on the exogenous variation only. That said, the results are very similar when running OLS regressions, suggesting that most of the variation in phase 1 experience is generated by our random assignment to different markets.

average.

Table 6: The impact of observational and personal payoff experience on responder minimal acceptable offers (MAO)

	Baseline and Role Switch Exp. only		Baseline and Full Information Exp. only		All Experiments	
	(1)	(2)	(3)	(4)	(5)	(6)
	MAO	MAO	MAO	MAO	MAO	MAO
Obs. experience	26.69*** (7.00)	15.74*** (5.31)	20.94 (14.80)	21.44* (11.60)	23.93 *** (6.57)	15.95 *** (5.11)
P. payoff experience	9.40 (7.16)	11.05 (6.82)	15.57 (10.19)	6.20 (9.03)	9.62 ** (4.33)	10.85 ** (5.36)
Constant	30.73*** (5.50)	38.76*** (3.97)	35.93*** (3.87)	36.08*** (3.51)	32.91 *** (3.05)	39.16 *** (3.19)
Adj. R^2	0.03	0.06	0.04	0.04	0.03	0.05
Observations	135	2025	162	2430	222	3330

Instrumental variables regressions estimating the impact of phase 1 personal payoff experience and observational experience on phase 2 MAO's. Estimates computed using the iterative GMM estimator. Columns (1) to (4) report results of pairwise combinations of the Baseline Experiment data with data from the Role Switch and the Full Information experiments, respectively. In column, personal payoff experience and observational experience are instrumented using 4 dummies, one for each phase 1 market treatment in each experiment. Columns (1) and (3) contain observations from period 1 of phase 2 only. Columns (2) and (4) contain data from all periods. Columns (5) and (6) combine data from all experiments. Here, personal payoff experience and observational experience are therefore instrumented using 6 dummies. Robust standard errors are clustered by phase 1 market matching group (2 clusters per session, 18 clusters in total in columns (1) and (2), 22 clusters in total in columns (3) and (4), 30 clusters in total in columns (5) and (6)). Significance levels: *** = 1%, ** = 5% and * = 10%. Columns (1) and (2) additionally contain dummies for high proposer personal payoff experience and high proposer observational experience. Columns (3) and (4) additionally contain a dummy for high proposer personal payoff experience (including proposer observational experience dummies would lead to collinearity, because high and low observational experience is unique to the Full Information experiment). Regressions in columns (5) and (6) additionally contain dummy variables for proposer personal payoff experience and for proposer observational experience.

Columns (1)-(4) of table 6 report the results, focusing on period 1 only in columns (1) and (3) and using data from all fifteen periods in columns (2) and (4). Although pooling data from only 2 out of 3 experiments doesn't always give us enough power to reach statistical significance at conventional thresholds, the regressions in the table are consistent in showing a large effect of both personal payoff experience and observational experience on responders' MAO's. What is also noteworthy is that the effects of past observational and payoff experience do not appear to differ much across the Role Switch experiment and the Full Information experiment. The impact of observational experience in the Role Switch experiment appears to be similar in magnitude to the effect of observational experience in the Full Information experiment.

Columns (5) and (6) of table 6 pool all three experiments for greater power and report reduced-form estimates of the impact of phase 1 observational experience and personal payoff experience on responder MAO's. Column (5) again focuses on period 1, while column (6) uses data from all 15 periods. In these pooled regressions, we again find that both observational experience and personal payoff experience have considerable effects, and this time we have enough power to reject the null hypothesis of no effect at $p < 0.05$ for personal payoff experience and at $p < 0.01$ for observational experience. Moreover, the estimates of the impact of observational experience and personal payoff experience are again very similar to the estimates in columns (1)-(4).

To get a rough sense of magnitudes, the regressions imply that a 10 percentage points increase in the average personal payoff experience increases first period minimum acceptable offers by approximately 1 chip. Similarly, a 10 percentage points increase in observed average offers increases first period minimum acceptable offers by another 2.4 chips. As expected, these effects are smaller when using data from all 15 periods, but they remain statistically significant.

Figure 6 shows how the effects of observational experience and personal payoff experience develop over the course of the 15 periods in phase 2 of the experiments. The left panel plots the coefficient estimate for observational experience, individually estimated for every period, similar to the regression for period 1 in column (5) of table 6. Dashed lines indicate 95% confidence intervals of the parameter estimates. Observational experience has a large and significant impact on minimal acceptable offers until period 8. In later periods, the effect gets smaller and is no longer significant at the 5% level. The right panel shows equivalent coefficient estimates for personal payoff experience. The effect of personal payoff experience remains relatively stable over the course of the 15 periods, with a slight downward trend.

3.4 Implications for Parameters of the Fairness Model

To quantitatively compare our framework with standard social preferences models that assume fixed fairness preferences, we now estimate the role that personal payoff experience and observational experience play in shaping either the reference point or the sensitivity parameters in equation (1). We do not consider our analysis here “structural” in the standard sense. Rather, we estimate linear models of how the fairness parameters— r , β , α —seem to depend on the history. In this sense, our analysis is analogous to the analysis in Section 3.3.

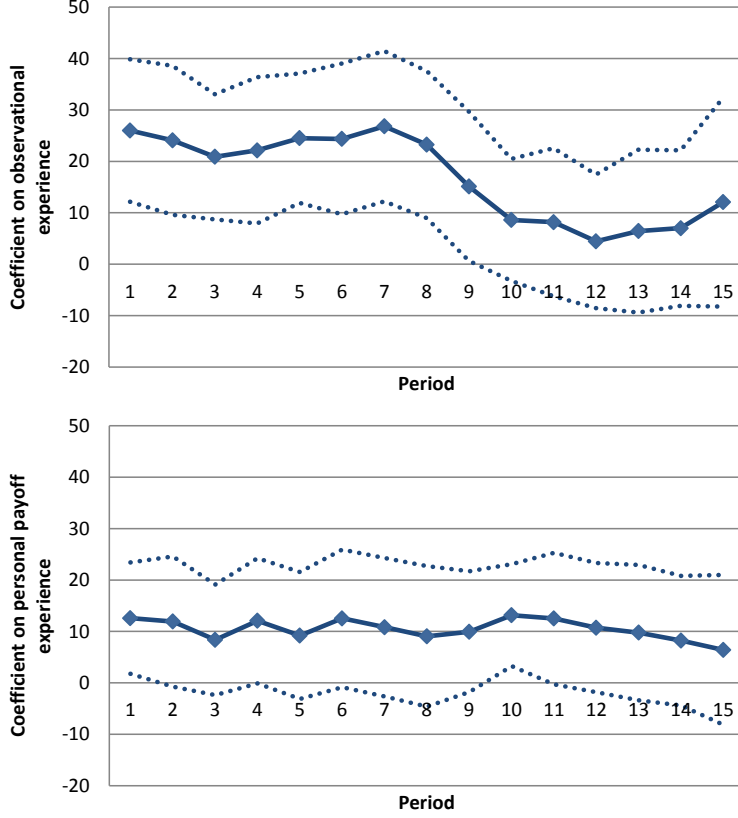


Figure 6: Top Panel: Development of the observational experience parameter over time, individually estimated for every period using the same IV regression specification as in column (5) of table 6. Bottom panel: Development of the personal payoff experience parameter over time, individually estimated for every period. Robust standard errors are clustered by the phase 1 market treatment group (2 clusters per session, 30 clusters in total). The dotted lines indicate 95% confidence intervals.

3.4.1 Fairness Reference Points

We model a responder's period τ reference point as being given by $r_R(\bar{\mu}_i^\tau, \bar{\nu}_i^\tau)$. Similarly, we model a proposer's period τ reference point as being given by $r_P(\bar{\mu}_i^\tau, 1 - \bar{\nu}_i^\tau)$.

To produce a tractable and estimable functional form, we assume that the reference point is formed through a convex combination of the average observational experience and the average personal payoff experience. Thus for responders,

$$r_R = (1 - \gamma_{PE} - \gamma_{OE})(1/2) + \gamma_{PE} \sum_{t=1}^{\tau-1} \mu^t / (\tau - 1) + \gamma_{OE} \sum_{t=1}^{\tau-1} \nu^t / (\tau - 1) \quad (4)$$

where γ_{PE} and γ_{OE} are the weights on personal payoff experience and observational experience. Similarly, for proposers,

$$r_P = (1 - \gamma_{PE} - \gamma_{OE})(1/2) + \gamma_{PE} \sum_{t=1}^{\tau-1} \mu^t / (\tau - 1) + \gamma_{OE} \sum_{t=1}^{\tau-1} (1 - \nu^t) / (\tau - 1) \quad (5)$$

The fairness sensitivity parameters are assumed to be exogenous, with $\beta = \beta_0$ and $\alpha = \alpha_0$ for constants α_0 and β_0 .

3.4.2 Sensitivity to (Un)Fairness

We also consider a model in which the fairness reference point is constant at the equal split norm— $r = 1/N$ in an N -player game, independent of past experiences—but the fairness sensitivity parameters β and α vary. We analogously define

$$\beta(\bar{\mu}_i^\tau, \bar{\nu}_i^\tau) = (1 - \gamma_{PE} - \gamma_{OE})(\beta_0) + \gamma_{PE} \sum_{t=1}^{\tau-1} \frac{\mu^t - m}{\tau - 1} + \gamma_{OE} \sum_{t=1}^{\tau-1} \frac{\nu^t - n}{\tau - 1} \quad (6)$$

$$= \beta_1 + \gamma_{PE} \sum_{t=1}^{\tau-1} \frac{\mu^t}{\tau - 1} + \gamma_{OE} \sum_{t=1}^{\tau-1} \frac{\nu^t}{\tau - 1} \quad (7)$$

$$\alpha(\bar{\mu}_i^\tau, \bar{\nu}_i^\tau) = (1 - \gamma_{PE} - \gamma_{OE})(\alpha_0) + \gamma_{PE} \sum_{t=1}^{\tau-1} \frac{\mu^t - m}{\tau - 1} + \gamma_{OE} \sum_{t=1}^{\tau-1} \frac{\nu^t - n}{\tau - 1} \quad (8)$$

$$= \alpha_1 + \gamma_{PE} \sum_{t=1}^{\tau-1} \frac{\mu^t}{\tau - 1} + \gamma_{OE} \sum_{t=1}^{\tau-1} \frac{\nu^t}{\tau - 1} \quad (9)$$

where $\beta_1 := (1 - \gamma_{PE} - \gamma_{OE})(\beta_0) - \gamma_{PE}m - \gamma_{OE}n$, $\alpha_1 := (1 - \gamma_{PE} - \gamma_{OE})(\alpha_0) - \gamma_{PE}m - \gamma_{OE}n$, and m and n are normalizing constants determining whether a particular observation or payoff experience increases or decreases the sensitivity to un(fairness). Our linear models for β and α thus involve constant terms β_1 and α_1 , as well as the slope parameters γ_{PE} and γ_{OE} that determine how *changes* in experience affect α and β . Although we cannot unpack the constant terms to separately identify α_0 and β_0 from m and n , our main interest is in γ_{PE} and γ_{OE} —how differences in observational and personal payoff experiences affect fairness sensitivity parameters.

We stress that our modeling of sensitivity to (un)fairness may not correspond perfectly to the true structural model, at least when interpreted literally. We show in Appendix C, however, that the salience model of Bordalo et al. (2012; 2013; 2015) generates fairness sensitivity weights that are monotonic in observed offers and experienced payoff shares. In the salience model, γ_{PE} and γ_{OE} correspond to the salience parameter multiplied by the fractions of consumers whose evoked sets contain individual payoff experience and observational payoff experience, respectively.

3.4.3 Estimation

As before, we use method of moments techniques to exploit purely exogenous variation created by random assignment to different treatments, and avoid potentially endogenous differences in MAOs in phase 1 and phase 2. We use equation (1) to compute a responder’s predicted MAO, given the predicted values of α , β , and r . The six different treatment cells (see table 4) allow us to use Generalized Method of Moments to recover γ_{PE} and γ_{OE} in equation (4) or in equation (6). The six moment conditions we obtain from the six different treatment cells are

$$E \left[\left(MAO - \frac{300\beta r}{2\beta r + \beta + 1} \right) T_j \right] = 0$$

where the T_j are dummies corresponding to the six possible treatment conditions. Letting $\xi = (\beta_0, \gamma_{PE}, \gamma_{OE})$ denote the parameters, the GMM estimator chooses the parameters $\hat{\xi}$ that minimize $(m(\xi) - \hat{m})'W(m(\xi) - \hat{m})$, where $m(\xi)$ are the theoretical moments, \hat{m} are the empirical moments, and W is the weighting matrix for the six moment conditions. The most efficient choice of W is the inverse of the variance-covariance matrix, which we approximate using an iterative estimation procedure as specified in Hall (2005). As always, we compute robust standard errors clustered at the phase 1 market matching group level. We use the Gauss-Newton algorithm to implement the minimum distance estimator.

A key difference from the empirical analysis in the rest of the paper is that because the fully specified model states that either the reference point r or the behindness aversion parameter β is shaped by *all* previous experience and observation, estimating our model on all 15 periods of phase 2 requires us to use data from phase 2 experiences and observations when formulating (4). Thus the past experience variables we construct use the average of all past experiences, including past experience from previous periods in phase 2. This is somewhat problematic because in phase 2, personal payoff experience and observational experience are highly collinear, as they are in experiment 1.³⁴ Thus when estimating our model on all 15 periods of phase 2, our only source of exogenous variation for separating between personal payoff experience and observational experience is still phase 1 experience. Because of the high degree of collinearity in phase 2, we are thus cautious about interpreting the estimates that arise from using all 15 periods of phase 2 data.

Table 7 presents the results, with columns (1) and (3) focusing on period 1 only, and columns (2) and (4) using data from all 15 periods. For the reasons mentioned above,

³⁴A related issue is that there is not a lot of exogenous variation in experience in phase 2. Most of the exogenous variation in phase 2 experience would have to come from exogenous variation in proposer offers, which comes from exogenous variation in proposers’ experience. In section 3.5 we confirm that the exogenous variation in proposers’ phase 1 experience does, indeed, impact their offers, at least initially.

Table 7: Parameter Estimates for the impact of personal payoff and observational experience in the fairness model

	Impact of γ_{PE} and γ_{OE} on the reference point r		Impact of γ_{PE} and γ_{OE} on the fairness sensitivity parameter β	
	(1)	(2)	(3)	(4)
γ_{PE}	0.13** (0.06)	0.04 (0.09)	0.37** (0.18)	0.10 (0.22)
γ_{OE}	0.33*** (0.08)	0.36*** (0.10)	0.97*** (0.28)	0.81*** (0.24)
λ	0.98*** (0.09)	0.78*** (0.08)		
Constant			0.31*** (0.08)	0.32*** (0.06)
Adj. R^2				
Observations	222	3330	222	3330
Hansen's J	1.65 (p=0.65)	0.30 (p=0.96)	1.16 (p=0.76)	0.20 (p=0.98)
Observations	222	3330	222	3330

Generalized Method of Moments estimates of the impact of observational and personal payoff experience on fairness reference points r and the sensitivity to unfairness β , using equations 1, 4 and 6. Columns (1) and (3) use period 1 data only, while columns (2) and (4) use data from all 15 periods. The six moment conditions are determined by instruments corresponding to the six Experiment \times Treatment conditions. Standard errors are clustered at the phase 1 market matching group level. The weighting matrix for the minimum distance estimator is computed iteratively to approximate the inverse of the variance-covariance matrix. The minimum distance estimation is implemented via the Gauss-Newton algorithm. Hansen's overidentification test reports the likelihood that the our specified model is consistent with the data. Significance levels: *** = 1%, ** = 5% and * = 10%.

the period 1 only data uses cleaner exogenous variation for estimating our three structural parameters. Column (1) shows that both personal payoff experience and observational experience receive positive weight in shaping the reference point when α and β are assumed to not be affected by phase 1 experiences. While both weights are significantly different from 0, the weight given to observational experience, 0.325, is roughly 2.5 times larger than the weight given to personal payoff experience, and this difference in weights is significant ($p = 0.07$). When estimating parameters using all data in column (2), the estimates look similar, with only γ_{PE} losing significance, potentially because of the partial collinearity problem.

Similarly, column (3) shows that when we assume that r is fixed exogenously at 0.5, personal payoff experience and observational experience both have a positive and significant impact on the sensitivity to unfairness β . Again, the relative impact of observational experience is roughly 2.6 times higher than the impact of personal payoff experience. This

difference, however, fails to be statistically significant at conventional levels ($p = 0.11$). When estimating the relative impact using all data in column (4), the coefficient on observational experience remains largely unchanged, while the coefficient on payoff experience becomes smaller and loses significance.

3.5 Proposer Offers

Last, we turn to proposer behavior, which should be influenced both by their own fairness motive as well as by responders' MAOs. To analyze the extent to which proposers' own fairness motives are path-dependent, we construct for each proposer his average personal payoff experience and observational experience from phase 1, and analyze how that influences his subsequent offer strategy. To analyze how proposers adjust their strategies to responder behavior, we use phase 1 exogenous variation in responders' experience: for each matching group we construct variables for the average phase 1 payoff experience and observational experience of all responders in the matching group. Table 8 displays regressions analyzing how proposers' and responders' phase 1 experience and observation influence proposers' phase 2 offers. To construct instruments for the 4 experience and observation variables we now use both variation in responder phase 1 conditions and in proposer phase 2 conditions. Because RC responders can be matched to either RC proposers or PC proposers (while PC responders are always matched to PC proposers), this gives us 3 different matching groups for each of the three experiments, for a total of $3 \times 3 = 9$ instruments. We use the iterative GMM estimator as before.

Consistent with the path-dependent fairness hypothesis, column (1) shows that in period 1, the estimated coefficient on personal payoff experience is negative and relatively large in magnitude, implying that proposers who are used to receiving a higher share of the surplus feel entitled to a greater share and thus are less likely to make a generous offer to responders. At the same time, the coefficient on observational experience is positive, implying that proposers who are used to observing higher offers are more likely to make a high offer.³⁵

Column (2) of the table, on the other hand, uses data from all 15 periods of phase 2. Over all 15 periods, proposers' offers are very sensitive to responders' experiences: The higher the payoff experience or observed offers of the responders, the higher the offers made by the proposers to these responders. In fact, the responder experience and observation coefficients in column (2) of Table 8 are almost identical to the responder experience and

³⁵As would be expected, the regression coefficients on average responder observational and payoff experience are insignificant. Because proposers were not aware of responders' history, there is no reason as to why proposers should react to responders' histories.

Table 8: The impact of observational and personal payoff experience on proposer offers

	(1)	(2)
	Offers	Offers
Observational experience	8.58*	7.58
	(4.50)	(4.86)
Personal payoff experience	-13.18***	5.88
	(2.59)	(3.69)
Avg. responder observational experience	5.94	15.98***
	(4.62)	(3.80)
Avg. responder personal payoff experience	-5.14	8.56**
	(4.43)	(3.40)
Constant	46.36***	39.36***
	(2.78)	(3.16)
Adj. R^2	0.02	0.11
Observations	222	3330

Instrumental variables regressions of offers on proposer and responder personal payoff experience and observational experience, using data from all three experiments. Both proposer and responder personal payoff experience and observational experience are instrumented using 9 dummies, one for each phase 2 matching group in each experiment. All estimates computed using the iterative GMM estimator. Columns (1) contains observations from period 1 of phase 2 only. Column (2) contains all observations. Standard errors are clustered at the phase 1 market matching group (2 clusters per session, 30 clusters in total). Significance levels: *** = 1%, ** = 5% and * = 10%.

observation coefficients in column (6) of Table 6, suggesting that proposers' offers respond almost one-for-one to responders' MAO's. The mechanism for this stark result is clear: proposers quickly learn what offers the responders find acceptable, and thus react quickly to match their behavior to that of responders. That proposers behave in a seemingly profit-maximizing manner here is not surprising, as they are less concerned with fairness under the model's assumption that $\alpha < \beta$.

4 Discussion

In this section we discuss what theories can and cannot explain our results. We suggest that "simple anchoring," commonly studied learning dynamics, and consumption-based models of dynamically-adjusting reference points are unlikely drivers of our results. We suggest that other theories such as backwards-looking fairness reference points or salience theory are more likely to be driving our results.

4.1 Reference Dependence

Backwards-looking Fairness Reference Points

One explanation of our results is a literal interpretation of our model of experience-based fairness reference points, formalized in Sections 2.2 and 3.4. This formulation follows Benjamin (2015), who proposes a variant of such a model and uses it to explain a number of puzzles in labor markets.

Expectations-based Fairness Reference Points

An alternative model, in the spirit of Kőszegi and Rabin (2006), is that the reference point is based on expectations.³⁶ A responder may choose to reject a proposer’s offer when that offer falls far short of what the responder expected to receive.

A model with *rational* expectations as reference points, however, would not be consistent with our main results. Once players learn which game they will be participating in for the subsequent 15 periods, their rational expectations about outcomes should not depend on their phase 1 experiences and observations.³⁷

Consequently, such a modification of our model would need to be combined with a model of *naïve*, rather than rational expectations in order to explain any path-dependence in responders’ phase 2 behavior. Adaptive expectations, combined with a theory of preferences in which responders like to reject offers that are below what they expected, could partly explain the impact of observational experience.³⁸ We leave it to future work to decompose

³⁶In the context of third-party punishment, Coffman (2010) tests the idea that third parties’ expectations may shape their punishment decisions, but does not find evidence for this hypothesis. In the context of risk preferences, Ericson et al. (2011); Heffetz and List (2014) and Camerer et al. (2016) provide (mixed) evidence that expectations shape reference points, which is also demonstrated by Abeler et al. (2011) and Camerer et al. (2016) in the context of effort provision.

³⁷Rational expectations should only be shaped by knowledge of the game structure, and beliefs about other players’ types. And since rational players should not have their beliefs systematically biased by play in different games, these rational players should not have different beliefs about each others’ types as a result of playing different games in phase 1. Of course, it may be possible to *accommodate* our results with a model in which there are multiple rational expectations equilibria and past experience serves as a coordination device for selecting an equilibrium. However, we do not find such an explanation particularly satisfactory, since it amounts to assuming a model with enough degrees of freedom in its predictions such that our data can’t falsify it. A more satisfactory account would have our empirical results as a *prediction*.

³⁸See Cooper and Dutcher (2011) for a sketch of a model based on belief-based reciprocity models (Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006). As we show in appendix B, however, such models do not naturally give rise to this effect. It is also interesting to note that such adaptive expectations should only be based on available information about other players’ behavior, which is captured by our observational experience variable. To explore this further, table 13 in Appendix H analyzes only the Full Information experiment, where all players receive the same information about phase 1 behavior. The table shows that the payoff differences generated by phase 1 market assignment still have a significant impact on behavior in this experiment. Consequently, further assumptions on the formation of naive expectations would need to be imposed to also explain the effect of personal payoff experience, such as overweighting of personally

how much of our observational experience effect is driven by backward-looking expectations.

Other Forms of Reference-dependent Preferences

Reference-dependent preferences have traditionally been discussed in the literature in the context of risky choice (Kahneman and Tversky, 1979), and have subsequently been applied to other domains of decision making, such as consumption (Kőszegi and Rabin (2006)). In these models, actual consumption, or monetary gains and losses, are evaluated relative to the reference point, and deviations of outcomes from this reference point are then associated with psychological gain/loss utility. In contrast, we posit that people’s fairness preferences, and not just consumption utility, is reference dependent. To what extent do these approaches differ? In Appendix B we show that models in which reference dependence affects only one’s utility from earnings, rather than perceptions of fairness, predict the *opposite* of our results in the Baseline experiment. Intuitively, the higher the payoffs in phase 1 of the experiment, the higher the reference point in phase 2, and thus the more painful it is to reject an offer and get a zero payoff.

4.2 Salience theory

Salience theory (Bordalo et al., 2012, 2013, 2015) provides one possible microfoundation for how the α and β parameters in our simple model might depend on past experience. We explore this in detail in Appendix C. We show that fairness will be more salient than own payoffs for responders who have previously observed and experienced high offers; and conversely for responders who have previously observed and experienced lower offers. We use a mixture model—in which some subjects’ “evoked sets” include observational experience and other subjects’ evoked sets include payoff experience—to rationalize the differential impact of payoff vs. observational experience on aggregate behavior.

The key feature of salience theory that makes it applicable to our results is that it delineates an important role for the effect of past experience in shaping salience weights. Other theories of context effects such as those of Kőszegi and Szeidl (2013) and Bushong et al. (2015), could not explain our results, as these theories assume that the focusing/attention weights are shaped solely by the decision-maker’s current choice set. Our empirical results about observational vs. personal payoff experience provide some new guidance on which elements of people’s histories are most likely to be part of their evoked sets.

experienced outcomes. Evidence that expectations may interact with fairness in other contexts is provided by Gilchrist et al. (2016), who find that unexpected bonuses promote more worker reciprocity than expected wage increases.

4.3 “Simple Anchoring”

Experimental evidence has shown that individuals can be influenced by arbitrary anchors (Lichtenstein and Slovic, 2006; Kahneman and Tversky, 2000; Ariely et al., 2003; Simonson and Tversky, 1992), and in a certain sense, the path-dependence effects that we identify could be related to anchoring. However, we are able to rule out a broad set of possible types of anchoring that we refer to as “simple anchoring.” We call “simple anchoring” the idea that a subject’s choice of action (e.g., offer) starts at some anchor, and then is incompletely adjusted toward the optimal choice of action. Formally, the choice of action is given by $a = (1 - \kappa)\vartheta + \kappa a^*$ where ϑ is the anchor, a^* is the optimal action, and $\kappa \in [0, 1]$ is the degree of adjustment away from the anchor.³⁹

Before moving on to a more thorough analysis, we first note that only the impact of observational experience can possibly be explained by simple anchoring, similar to the findings by Ariely et al. (2003). Observational experience captures all observational attributes of phase 1, such as high and low offers, that could serve as arbitrary anchors. The significant effect of personal payoff experience cannot be explained by such simple anchoring. A second immediate argument against such simple forms of anchoring is the change in strategy space between phase 1 (accept/reject) and phase 2 (MAO), which further limits the applicability of “simple anchoring”.

To further distinguish simple anchoring from other theories, in Appendix A we show theoretically that *differences in preferences* will lead to differences in behavior in environments such as the Ultimatum Game, but that they will not lead to differences in behavior in competitive market games as in phase 1 of our experiment (see Appendix A). In contrast, simple anchoring predicts that there should be differences in behavior in both phase 2 *and phase 1*. Building on this, in Appendix G we provide evidence that observational experience is also unlikely to be the consequence of simple anchoring. We do this using the Full Information experiment, in which all subjects received feedback about the average offers in both the PC and the RC markets after every period during phase 1 of the experiment. Thus subjects in the PC market observe a potential anchor that is substantially lower than offers in the PC market, whereas subjects in the RC market observe a potential anchor that is substantially higher than offers in the RC markets. As we have already demonstrated, such observational experience has large effects on behavior in phase 2. But if subjects’ behavior were indeed driven by simple anchoring, then responders in the RC market in the Full Information experiment should show higher acceptance rates than responders in the RC

³⁹For example, Ariely et al. (2003) have shown that the provision of arbitrary anchors, such as the final two digits of one’s social security number, affect an individual’s willingness to pay.

market in the Baseline or in the Role Switch experiments.⁴⁰ Similarly, PC Proposers in the Full Information experiment should offer less than PC Proposers in the other two experiments, and RC proposers in the Full Information experiment offer more than RC proposers in the other two experiments. Contrary to this, both proposers and responders in the Full Information experiment behave identically to the proposers and responders in the Baseline and the Role Switch experiments.

4.4 Reinforcement Learning

Roth and Erev (1995)⁴¹ have applied the reinforcement learning (RL) model to ultimatum game dynamics. There are several aspects of our data that are inconsistent with reinforcement learning models. Plainly, our results about the effects of observational experience are inconsistent with RL, as RL is based only on experience, rather than observation. Relatedly, applying the RL model to our experiments is importantly limited by the fact that the strategy spaces change significantly between phases 1 and 2 in each of our experiments.⁴² The RL model is a model that is applied to learning over time with a fixed strategy space.⁴³

4.5 Reciprocity vs. Distributional Preferences

Literally interpreted, our intentionally simple theoretical framework is a model of path-dependent *distributional* preferences. Alternatively, we could base our analysis on the Charness and Rabin (2002) model, in which responders want to punish proposers who don't take "nice" actions. The corresponding generalization of this model would similarly posit that past experience could either modify the reference point for what is considered "nice," or it could modify responders' sensitivity to not being treated nicely. Our model could be interpreted as a reduced-form manifestation of this kind of framework, where r is the fairness reference point and β is the strength of (negative) reciprocity. We cannot distinguish between distributional preferences and reciprocity preferences in our data, and we leave it for future work to determine whether past experiences shape distributional preferences, reciprocity

⁴⁰We restrict attention to the RC market because in the PC market, responders almost never reject both offers, and hence there is not enough variance in the data to identify a potential impact.

⁴¹See also Grosskopf (2003)

⁴²In the Baseline and Full Information experiments, the strategy space is {Accept, Reject} in Phase 1, while in Phase 2 the strategy space is {0, 5, ..., 100}. Perhaps even more problematic is the Role Switch experiment, in which a subject chooses between offers in Phase 1, but chooses MAOs in Phase 2.

⁴³Our result that proposers quickly adjust their offers to match the behavior of responders is consistent with RL. However, it is also explained by our preference-based model, which specifies that $\beta > \alpha$; e.g., that subjects care more about being behind than being ahead. Indeed, we show in Appendix C that if $\alpha = 0$ then proposers would essentially behave as pure profit-maximizers and thus their behavior would only be affected by responder experience.

preferences, or both.

5 Conclusion

While most work on social preferences has progressed under the presumption of static preferences, we show that fairness preferences are malleable and shaped by economic forces that lead to the experience of different types of market outcomes. We also show that such malleability can be captured by simple, tractable, and estimable models. Our reduced-form and structural results imply significant deviations from existing models of fairness. Our results therefore highlight the importance of considering contrast effects in the domain of fairness. In addition to the labor market implications explored by Benjamin (2015) and others, our evidence of path dependence also has implications for various settings studied in industrial organization. Our results imply that a tradeoff exists between the immediate loss of customers who judge a certain transaction to be unfair, and the long run profits generated through an increased willingness to pay of customers once the fairness preferences have adjusted. This leads to new considerations for dynamic price-setting.

Our results on observational experience also generate practical implications for price discrimination. A firm trying to price discriminate among consumers should try to conceal this price discrimination from consumers being offered the highest price, but inform those consumers receiving low prices. When such differential information provision is not feasible, the formal models we have introduced could be used to analyze when shrouding price discrimination is payoff maximizing. With the rise of the potential for personalized pricing and advertising in e-commerce, such considerations become increasingly important.

More generally, a key implication of our results on observational experience is that informing buyers (workers) about other prices (wages) should change the prices (wages) that are perceived as acceptable, even when such information is payoff irrelevant. Consequently, increased information dispersion should have the effect of homogenizing fairness norms.

A number of questions about mechanisms remain unanswered. For example, further work is needed to precisely identify which theories explain the path-dependence that we identify. Our experimental results also don't shed light on the precise nature of how past experiences are coded: Do more recent experiences receive more weight? Are past offers that are below, rather than above, what one is used to particularly salient? These and other theoretical and empirical extensions of our analysis are directions for future research.

References

- Abeler, Johannes, Armin Falk, Lorenz Goette, and David Huffman, "Reference points and effort provision," *The American Economic Review*, 2011, pp. 470–492.
- Akerlof, George A, "Labor contracts as partial gift exchange," *The Quarterly Journal of Economics*, 1982, *97* (4), 543–569.
- Ariely, Dan, George Loewenstein, and Drazen Prelec, "'Coherent arbitrariness': Stable demand curves without stable preferences," *The Quarterly Journal of Economics*, 2003, *118* (1), 73–106.
- Bartling, Bjoern, Roberto A. Weber, and Lan Yao, "Do Markets Erode Social Responsibility," *The Quarterly Journal of Economics*, 2015, *130* (1), 219–266.
- Bartling, Björn, Manuel Grieder, and Christian Zehnder, "Does Competition Justify Inequality," CESifo Working Paper Series 5651, CESifo Group Munich 2015.
- Bauer, Michal, Alessandra Cassar, Julie Chytilová, and Joseph Henrich, "War's enduring effects on the development of egalitarian motivations and in-group biases," *Psychological science*, 2014, *25* (1), 47–57.
- Bednar, Jenna, Yan Chen, Tracy Xiao Liu, and Scott Page, "Behavioral spillovers and cognitive load in multiple games: An experimental study," *Games and Economic Behavior*, 2012, *74* (1), 12–31.
- Benabou, Roland, Jean Tirole et al., "Laws and Norms," Technical Report 6290, Institute for the Study of Labor (IZA) 2012.
- Benjamin, Daniel J., "A Theory of Fairness in Labor Markets," *Japanese Economic Review*, 2015, *66* (2), 182–225.
- Bettinger, Eric and Robert Slonim, "Using experimental economics to measure the effects of a natural educational experiment on altruism," *Journal of Public Economics*, 2006, *90* (8), 1625–1648.
- Bhargava, Saurabh and Ray Fisman, "Contrast effects in sequential decisions: Evidence from speed dating," *Review of Economics and Statistics*, 2014, *96* (3), 444–457.
- Binmore, Ken, Peter Morgan, Avner Shaked, and John Sutton, "Do people exploit their bargaining power? An experimental study," *Games and Economic Behavior*, 1991, *3* (3), 295–322.

- Bolton, Gary E. and Axel Ockenfels**, “ERC: A Theory of Equity, Reciprocity, and Competition,” *American Economic Review*, 2000, *90* (1), 166–93.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Salience Theory of Choice Under Risk,” *Quarterly Journal of Economics*, 2012, pp. 1243–1285.
- , – , and – , “Salience and Consumer Choice,” *Journal of Political Economy*, 2013, *121* (5), 803–843.
- , – , and – , “Memory, attention and choice,” Technical Report 2015.
- Bowman, David, Deborah Minehart, and Matthew Rabin**, “Loss aversion in a consumption–savings model,” *Journal of Economic Behavior & Organization*, 1999, *38* (2), 155–178.
- Brandts, Jordi, Matthew Ellman, and Gary Charness**, “Let’s talk: How communication affects contract design,” *Journal of the European Economic Association*, forthcoming.
- Bushong, Benjamin, Mathew Rabin, and Josh Schwartzstein**, “A Model of Relative Thinking,” 2015, *working paper*.
- Camerer, Colin and Teck Hua Ho**, “Experience-weighted Attraction Learning in Normal Form Games,” *Econometrica*, 1999, *67* (4), 827–874.
- Camerer, Colin F, Anna Dreber, Eskil Forsell, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Johan Almenberg, Adam Altmejd, Taizan Chan et al.**, “Evaluating replicability of laboratory experiments in economics,” *Science*, 2016, *351* (6277).
- Cason, Timothy N, Anya C Savikhin, and Roman M Sheremeta**, “Behavioral spillovers in coordination games,” *European Economic Review*, 2011, *56*, 233–245.
- Cassar, Alessandra, Pauline Grosjean, and Sam Whitt**, “Legacies of violence: trust and market development,” *Journal of Economic Growth*, 2012, pp. 1–34.
- Charness, Gary and Matthew Rabin**, “Understanding Social Preferences with Simple Tests,” *Quarterly Journal of Economics*, 2002, *117* (3), 817–869.
- Coffman, Lucas C.**, “Essays in Experimental Economics,” *Doctoral Dissertation*, 2010.
- Cooper, David J. and E. Glenn Dutcher**, “The dynamics of responder behavior in ultimatum games: a meta-study,” *Experimental Economics*, 2011, *14*, 519–546.

- Dufwenberg, Martin and Georg Kirchsteiger**, “A Theory of Sequential Reciprocity,” *Games and Economic Behavior*, 2004, 47 (2), 268–298.
- Ericson, Keith M Marzilli, Andreas Fuster et al.**, “Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments,” *The Quarterly Journal of Economics*, 2011, 126 (4), 1879–1907.
- Falk, Armin and Nora Szech**, “Morals and markets,” *science*, 2013, 340 (6133), 707–711.
- **and Urs Fischbacher**, “A theory of reciprocity,” *Games and Economic Behavior*, 2006, 54 (2), 293–315.
- **, Ernst Fehr, and Christian Zehnder**, “Fairness perceptions and reservation wages—the behavioral effects of minimum wage laws,” *The Quarterly Journal of Economics*, 2006, 121 (4), 1347–1381.
- Fehr, Ernst and Klaus M. Schmidt**, “A Theory of Fairness, Competition, and Cooperation,” *The Quarterly Journal of Economics*, 1999, 114 (3), 817–868.
- **, Georg Kirchsteiger, and Arno Riedl**, “Does fairness prevent market clearing? An experimental investigation,” *The Quarterly Journal of Economics*, 1993, 108 (2), 437–459.
- **, Lorenz Goette, and Christian Zehnder**, “A Behavioral Account of the Labor Market: The Role of Fairness Concerns,” *Annual Review of Economics*, 2009, 1 (1), 355–384.
- **, Oliver D. Hart, and Christian Zehnder**, “How Do Informal Agreements and Renegotiation Shape Contractual Reference Points?,” *Journal of the European Economic Association*, 2015, 13 (1), 1–28.
- **, Oliver Hart, and Christian Zehnder**, “Contracts as Reference Points—Experimental Evidence,” *American Economic Review*, 2011, 101 (2), 493–525.
- Fischbacher, Urs**, “z-Tree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics*, 2007, 10, 171–178.
- **, Christina M. Fong, and Ernst Fehr**, “Fairness, errors and the power of competition,” *Journal of Economic Behavior & Organization*, 2009, 72 (1), 527–545.
- Frederick, Shane**, “Cognitive reflection and decision making,” *The Journal of Economic Perspectives*, 2005, 19 (4), 25–42.

- Fudenberg, Drew, David K. Levine, and Zacharias Maniadis**, “On the Robustness of Anchoring Effects in WTP and WTA Experiments,” *American Economic Journal: Microeconomics*, 2012, 4 (2), 131–145.
- Gilchrist, Duncan S, Michael Luca, and Deepak Malhotra**, “When $3+1 > 4$: Gift structure and reciprocity in the field,” *Management Science*, 2016, 62 (9), 2639–2650.
- Gilligan, Michael J., Benjamin J. Pasquale, and Cyrus Samii**, “Civil War and Social Cohesion: Lab-in-the-Field Evidence from Nepal,” *American Journal of Political Science*, 2014, 58 (3), 604–619.
- Greiner, Ben**, “The Online Recruitment System ORSEE: A Guide for the Organization of Experiments in Economics,” Discussion Papers on Strategic Interaction 2003–10, Max Planck Institute of Economics, Strategic Interaction Group 2004.
- Grimm, Veronika and Friederike Mengel**, “An experiment on learning in a multiple games environment,” *Journal of Economic Theory*, 2012, 147, 2220–2259.
- Grosskopf, Brit**, “Reinforcement and directional learning in the ultimatum game with responder competition,” *Experimental Economics*, 2003, 6 (2), 141–158.
- Güth, Werner, Rolf Schmittberger, and Bernd Schwarze**, “An experimental analysis of ultimatum bargaining,” *Journal of Economic Behavior & Organization*, 1982, 3 (4), 367–388.
- Hall, Robert E**, “Employment fluctuations with equilibrium wage stickiness,” *American economic review*, March 2005, 95 (1), 50–65.
- Handgraaf, Michel JJ, Eric Van Dijk, and David De Cremer**, “Social utility in ultimatum bargaining,” *Social Justice Research*, 2003, 16 (3), 263–283.
- Hargreaves-Heap, Shaun and Yanis Varoufakis**, “Some Experimental Evidence on the Evolution of Discrimination, Co-operation and Perceptions of Fairness,” *The Economic Journal*, 2002, 112 (481), 679–703.
- Hart, Oliver**, “Hold-up, Asset Ownership, and Referene Points,” *The Quarterly Journal of Economics*, 2009, 124 (1), 267–200.
- **and Bengt Holmstrom**, “A Theory of Firm Scope,” *The Quarterly Journal of Economics*, 2010, 125 (2), 483–513.

- **and John Moore**, “Contracts as Reference Points,” *The Quarterly Journal of Economics*, 2008, *123* (1), 1119–11158.
- Hartzmark, Samuel M and Kelly Shue**, “A tough act to follow: Contrast effects in financial markets,” *Mimeo*, 2015.
- Heffetz, Ori and John A List**, “Is the endowment effect an expectations effect?,” *Journal of the European Economic Association*, 2014, *12* (5), 1396–1422.
- Herz, Holger and Dmitry Taubinsky**, “Market experience is a reference point in judgments of fairness,” Technical Report 128, Working Paper Series, University of Zurich, Department of Economics 2013.
- Kagel, John H, Chung Kim, and Donald Moser**, “Fairness in ultimatum games with asymmetric information and asymmetric payoffs,” *Games and Economic Behavior*, 1996, *13* (1), 100–110.
- Kahneman, Daniel and Amos Tversky**, “Prospect theory: An analysis of decision under risk,” *Econometrica: Journal of the Econometric Society*, 1979, pp. 263–291.
- **and** –, *Choices, Values, and Frames*, 1st ed., New York: Cambridge University Press, 2000.
- , **Jack L Knetsch, and Richard Thaler**, “Fairness as a constraint on profit seeking: Entitlements in the market,” *The American economic review*, September 1986, *76* (4), 728–741.
- Kaur, Supreet**, “Nominal Wage Rigidity in Village Labor Markets,” *Working Paper*, 2012.
- Kenrick, Douglas T and Sara E Gutierrez**, “Contrast Effects and Judgments of Physical Attractiveness: When Beauty Becomes a Social Problem,” *Journal of Personality and Social Psychology*, 1980, *38* (1), 131–140.
- Kőszegi, Botond and Adam Szeidl**, “A Model of Focusing in Economic Choice,” *Quarterly Journal of Economics*, 2013, *128* (1), 53–104.
- **and Matthew Rabin**, “A Model of Reference-Dependent Preferences.,” *Quarterly Journal of Economics*, 2006, *121* (4), 1133–65.
- Lévy-Garboua, Louis, David Masclet, and Claude Montmarquette**, “A behavioral Laffer curve: Emergence of a social norm of fairness in a real effort experiment,” *Journal of Economic Psychology*, 2009, *30* (2), 147–161.

- Lichtenstein, Sarah and Paul Slovic**, *The construction of preference*, 1st ed., New York: Cambridge University Press, 2006.
- List, John A., Zacharias Maniadis, and Fabio Tufano**, “One Swallow Doesn’t Make a Summer: New Evidence on Anchoring Effects,” *American Economic Review*, *forthcoming*, 2013.
- Pepitone, Albert and Mark DiNubile**, “Contrast effects in judgments of crime severity and the punishment of criminal violators,” *Journal of Personality and Social Psychology*, 1976, *33* (4), 448–459.
- Peysakhovich, Alexander and David G. Rand**, “Habits of Virtue: Creating Norms of Cooperation and Defection in the Laboratory,” *Management Science*, 2015, *62* (3), 631–647.
- Rabin, Matthew**, “Incorporating Fairness into Game Theory and Economics,” *American Economic Review*, 1993, *83* (5), 1281–1302.
- Rao, Gautam**, “Familiarity Does not Breed Contempt: Diversity, Discrimination, and Generosity in Delhi Schools,” *working paper*, 2013.
- Roth, Alvin E and Ido Erev**, “Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term,” *Games and economic behavior*, 1995, *8* (1), 164–212.
- , **Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir**, “Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study,” *The American Economic Review*, December 1991, *81* (5), 1068–1095.
- Silverman, Dan, Joel Slemrod, and Neslihan Uler**, “Distinguishing the Role of Authority ”In” and Authority ”To”,” *Journal of Public Economics*, 2014, *104* (6), 1518–1550.
- Simonson, Itamar and Amos Tversky**, “Choice in context: tradeoff contrast and extremeness aversion.,” *Journal of Marketing Research*, 1992.
- Skott, Peter**, “Fairness as a source of hysteresis in employment and relative wages,” *Journal of Economic Behavior & Organization*, 2005, *57* (3), 305–331.
- Voors, Maarten J, Eleonora EM Nillesen, Philip Verwimp, Erwin H Bulte, Robert Lensink, and Daan P Van Soest**, “Violent conflict and behavior: a field experiment in Burundi,” *The American Economic Review*, 2012, *102* (2), 941–964.

Zwick, Rami and Vincent Mak, “Gaming with Fairness: Some Conjectures on Behavior in Alternating Offer Bargaining Experiments,” in Gary Bolton and Rachel Croson, eds., *Conflict Resolution Handbook*, The Oxford University Press, 2012, pp. 91–107.

Appendix

A Results for Market Games

In this appendix we consider predictions for the PC and RC market games. Proposers' preferences (β_P, r_P) are drawn from some distribution F_P and responders' preferences (β_R, r_R) are drawn from some distribution F_R . The solution concept we consider is a pure strategy Perfect Bayesian Equilibrium.

1. *There is a unique PBE in the PC market game: both proposers offer 100 chips, and the responder randomly chooses one of the offers.*

Proof. [sketch] Let \underline{a} be Proposer 1's offer, and suppose that $\underline{a} < 100$. Clearly, it is never strictly optimal for Proposer 2 to offer anything less than \underline{a} . It is also not optimal to offer \underline{a} . By choosing $\underline{a} + \epsilon$ instead of \underline{a} , the proposer increases his chance of winning from $1/2$ to 1 . The net impact on utility is $(100 - \underline{a} - \epsilon) - (100 - \underline{a})/2 > 0$ for ϵ sufficiently small. Thus, it is not possible to have $\underline{a} < 100$ in equilibrium. □

2. *The following is a PBE of the RC market game: the proposer offers 0 chips and the responders accept with probability 1.*

Proof. [sketch] First, we show that the Proposer offering 0 and the responders accepting all offers is an equilibrium. Since $\alpha < 1$, it is always optimal for a proposer to offer nothing if that offer will be accepted with probability 1. Next consider a responder's best response function in this proposed equilibrium. Note that by deviating and choosing to reject the offer, the responder cannot decrease the size of the pie. Since the other responder will accept, the final outcome after the deviation will still be the proposer getting 100 chips and the responders getting nothing. Thus there is no incentive to deviate. □

B Appendix to Section 4

B.1 Reference dependence over monetary payoffs

Here we consider a model in which preferences are reference-dependent with respect to monetary payoffs, but the fairness preferences are fixed. Let a responders' preferences be given by $u(\pi_R, \pi_P | r) = \pi_R + \mu(\pi_R - r) + f(\pi_R, \pi_P)$ where r is the reference point, μ is the gain-loss utility, and f is the fairness utility. Let μ and f be continuous and assume that

$u(\pi_R, 300 - 3\pi_R|r)$ is strictly increasing in π_R for each r . This guarantees that for each value of r , there is a minimally acceptable offer $MAO(r)$. Finally, we make the standard assumption that μ is concave.⁴⁴

3. $MAO(r)$ is decreasing in r .

The intuition for the Proposition is simple: the higher the reference point r , the more painful it is to get a payoff of 0, and thus the lower the MAO. Formally, $\mu(\pi_R - r) - \mu(0 - r)$ is decreasing in r because of the concavity of μ .

Now the natural assumption to make about how experience shapes the reference point is that experiencing higher offers or payoffs should lead to a higher reference point r . This is the case in theories of backwards-looking reference points, as in Bowman et al. (1999). Alternatively, if r is shaped by expectations which are based on the types of offers observed, then r should again be increasing in history of offers. Thus in experiment 1, r should be higher for responders in the PC market than for responders in the RC market. The Proposition thus implies that PC responders should actually have lower MAO's than RC responders - the opposite of our experimental results. This implies that unless fairness preferences are themselves shaped by past experience, there is no natural model of reference-dependence over monetary payoffs that is consistent with our experimental results.

B.2 Belief-based reciprocity models

We now argue that belief-based reciprocity models do not offer a natural explanation of our results. First, consider Dufwenberg and Kirchsteiger's (2004) extension of Rabin's (1993) intention based reciprocity model. For the proposer, $[0, 100]$ is the efficient set of offers a , and thus the "equitable" benchmark is given by $a^{ep} = 50$. If the proposer believes that a responder accepts his offer of a with probability θ , then his kindness toward this responder is thus $\theta a - 50$.

For the responder, the set of efficient strategies conditional on an offer a is simply to accept; thus the "equitable" payoff to the proposer conditional on an offer a is simply a . Thus if a responder rejects a proposer's offer, his kindness toward the proposer is $-a$. Letting ϕ denote the strength of the reciprocity motive, the responder's payoff from choosing an MAO is given by

$$\int_x [a - \phi a(50 - a\tilde{\theta})\mathbf{1}_{x \geq M}] d\tilde{F}(x) \tag{10}$$

⁴⁴Note that our assumptions on μ are more general than those of Köszegi and Rabin (2006).

where \tilde{F} are the responder's beliefs about the strategy of the proposer, and $\tilde{\theta}$ is the probability that the proposer thinks his offer will be accepted (from the responder's perspective; i.e., it is the responder's second order belief). Thus the MAO is given simply as the value M for which $M - \phi M(50 - M\tilde{\theta}) = 0$. The key feature of this condition is that M does not depend on the responder's expectation of proposer behavior. Thus this model, combined with a reasonable theory of adaptively formed expectations, would still not explain our results.

Similarly, in Falk and Fischbacher's (2006) model, a straightforward extension of (19) in their Appendix 3, shows that $U_{2A} - U_{2R}$, the relative gain from accepting versus rejecting an offer a , is given by

$$U_{2A} - U_{2R} = a + \phi\tilde{\theta}(300 - 3a - a)(100 - a) \quad (11)$$

Again, this shows that the acceptance decision does not depend on beliefs \tilde{F} .

It is possible that past experiences might shape the second order belief $\tilde{\theta}$. However, there are no models of learning that posit how past experiences shape second-order beliefs, and there isn't a clear hypothesis about how our phase 1 experiences should shape it. Moreover, intention based reciprocity models do not make a robust prediction about how $\tilde{\theta}$ affects M . In Dufwenberg and Kirchsteiger (2004), M is decreasing in $\tilde{\theta}$ while in Falk and Fischbacher (2006) M is increasing in $\tilde{\theta}$ (note that the utility from accepting is decreasing in $\tilde{\theta}$ in equation (11) while it is increasing in equation (10)).

C Salience Theory

C.1 Set up

The true utility function is $U_i = \Pi [\mu_i - \beta \max(r - \mu_i, 0) - \alpha \max(\mu_i - r, 0)]$, where $\mu_i = \pi_i/\Pi$ is player i 's share of the pie, and $r = 1/N$ is the fairness norm in an N -player game.

How a responder trades off payoffs and fairness, however, depends on which is more salient. We let $\Delta_i := \mu_i - r$ denote how much the responder's share of the pie deviates from the fairness norm, and we let (μ_i, Δ_i) denote the pair consisting of the responder's payoff and the fairness deviation. The salience of payoffs versus fairness deviations depends on how those attributes depart from payoff and fairness values in the responder's evoked set. The evoked set \mathcal{E} includes 1) the option corresponding to the proposer's offer, $(\mu(a_i), \Delta(a_i))$, where $\mu(a_i) = \frac{a_i}{3(100-a_i)+a_i} = \frac{a_i}{300-2a_i}$ and $\Delta(a_i) := \frac{a_i}{3(100-a_i)+a_i} - r = \frac{a_i}{300-2a_i} - r$. It includes 2) the option corresponding to the responder rejecting, $(0, 0)$. And it includes 3) a historical average of payoffs and fairness values either experienced or observed by the responder, (μ_i^H, Δ_i^H) . We discuss the construction of the historical averages later.

Given the evoked set \mathcal{E} , the salience is determined by comparison to the reference good $(\bar{\mu}, \bar{\Delta})$ given by $\bar{\mu} = \frac{\mu_i + 0 + \mu_i^H}{3}$ and $\bar{\Delta} = \frac{\Delta(\mu_i) + 0 + \Delta^H}{3}$. The salience of a payoff π_i is given by $\sigma(\mu_i, \bar{\mu})$ and the salience of a fairness deviation Δ_i is given by $\sigma(\Delta_i, \bar{\Delta})$. As in BGS, we assume that the salience function $\sigma(\cdot, \cdot)$ satisfies

1. Scale invariance. $\sigma(\alpha x, \alpha y) = \sigma(x, y)$
2. Diminishing sensitivity. If $x \geq 0$, then for any $\epsilon > 0$, $\sigma(x + \epsilon, y + \epsilon) \leq \sigma(x, y)$, with the inequality strict if $y > 0$.
3. Reflection. $\sigma(x, y) = \sigma(-x, -y)$.

A salience function that satisfies these properties for $x, y \neq 0$ is $\sigma(x, y) = \frac{|x-y|}{|x|+|y|}$. Note that the maximum value of this salience function is 1, and that $\sigma(x, 0) = \sigma(0, y) = 1$ for $x, y \neq 0$. To define $\sigma(0, 0)$, we thus make the more general fourth assumption that $\sigma(0, y) = \sigma(x, 0) = 1$ for all x, y .

As in BGS, the responder evaluates the option (μ_i, Δ_i) as follows:

$$U_i^s = \begin{cases} \Pi [\mu_i - \delta (\beta \max(r - \mu_i, 0) + \alpha \max(\mu_i - r, 0))] & \text{if } \sigma(\mu_i, \bar{\mu}) > \sigma(\Delta_i, \bar{\Delta}) \\ \Pi [\mu_i - (\beta \max(r - \mu_i, 0) + \alpha \max(\mu_i - r, 0))] & \text{if } \sigma(\mu_i, \bar{\mu}) = \sigma(\Delta_i, \bar{\Delta}) \\ \Pi [\delta \mu_i - (\beta \max(r - \mu_i, 0) + \alpha \max(\mu_i - r, 0))] & \text{if } \sigma(\mu_i, \bar{\mu}) < \sigma(\Delta_i, \bar{\Delta}) \end{cases}$$

where $\delta \in (0, 1]$. To ensure that the responder does not reject offers that give him more than half of the pie, we make the reasonable assumption that $\alpha/\delta < 2$.

C.2 Explaining Experiment 1 results

The perceived utility of rejection is always 0. How does the responder evaluate the utility of accepting the action? We first begin with the case in which $(\mu_i^H, \Delta_i^H) \in \{(0, -1/3), (1, 2/3)\}$. These two cases correspond to the equilibrium outcomes in the PC and RC markets, respectively.

Case 1. $(\mu_i^H, \Delta_i^H) = (1, 2/3)$ In this case, $\bar{\mu} = (1 + \mu_i)/3$ and $\bar{\Delta} = ((\mu_i - 1/2) + 2/3)/3$. Thus payoff salience is given by $\sigma(\mu_i, (1 + \mu_i)/3) = \sigma(3\mu_i, \mu_i + 1)$, while fairness salience is given by $\sigma(\mu_i - 1/2, (\mu_i + 1/6)/3) = \sigma(3\mu_i - 3/2, \mu_i + 1/6)$. Now because the proposer makes an offer that gives the responder a payoff smaller than his own, $\mu_i < 1/2$, and thus the combination of the scale invariance and diminishing sensitivity assumptions implies that

fairness is always more salient than payoffs in this case. The smallest share μ_i a responder is willing to accept in this case must satisfy

$$\delta\mu_i - \beta(1/2 - \mu_i) = 0$$

and thus $\mu_i = \frac{\beta}{2(\beta+\delta)}$. Since $\mu_i = \frac{a_i}{300-2a_i}$ by definition, this shows that smallest offer a responder is willing to accept in this case is thus $M = \frac{150\beta}{\delta+2\beta}$.

Case 2. $(\mu_i^H, \Delta_i) = (0, -1/3)$ In this case, $\bar{\mu} = \mu_i/3$ and $\bar{\Delta} = (\mu_i - 5/6)/3$. Thus payoff salience is given by $\sigma(\mu_i, \mu_i/3) = \sigma(3, 1)$, while fairness salience is given by $\sigma(\mu_i - 1/2, (\mu_i - 5/6)/3) = \sigma(3\mu_i - 3/2, \mu_i - 5/6)$. Now because the proposer makes an offer that gives the responder a payoff smaller than his own, $\mu_i < 1/2$ and thus Δ_i and $\bar{\Delta}$ are negative. In this case, payoffs will be more salient than payoffs if and only if $\frac{\mu_i - 5/6}{\mu_i - 3/2} < 3$, which happens when $\mu_i < 11/24$. When payoffs are salient, the smallest acceptable share is the solution to

$$\mu_i - \delta\beta(1/2 - \mu_i) = 0$$

and thus is $\mu_i = \frac{\delta\beta}{2(1+\delta\beta)}$. Now since $\mu_i = \frac{a_i}{300-2a_i}$, the smallest offer that a responder is willing to accept in this case is $M_1 = \frac{150\delta\beta}{1+2\delta\beta}$. When fairness is salient, the smallest acceptable offer is given by $M_2 = \frac{150\beta}{\delta+2\beta}$ as before. Because $M_1 < M_2$, it follows that the proposers will offer M_1 iff $\frac{\delta\beta}{2(1+\delta\beta)} \geq 11/24$ or, equivalently, iff $\delta\beta < 11$. If, however, $\delta\beta \geq 11$ then proposers offer M_2 . Note, however, that $\delta\beta < 11$ is a very general condition that is extremely unlikely to be violated. In the standard ultimatum game without salience considerations, a $\beta > 11$ would imply that responders reject all offers that give them less than 42.3% of the total pie.

Proposition 1. *Suppose that $\delta\beta \geq 11$. Then $MAO_{RC} = \left(\frac{\delta^2+2\delta\beta}{1+2\delta\beta}\right) MAO_{PC} < MAO_{PC}$.*

C.3 The Mixture Model and the Role Switch and Full Information experiments

Is it observational experience or personal payoff experience that enters into the evoked set \mathcal{E} as (μ^H, Δ^H) ? While this does not affect our experiment 1 interpretation, since the two are essentially identical, this can affect behavior in the Role Switch and the Full Information experiments. We suppose that a fraction q relies on observational experience, and a fraction $1 - q$ relies on personal payoff experience. Out of those who rely on observational experience, we suppose that in the Full Information Experiment, a fraction ω relies on what they observe in their own market, while a fraction $(1 - \omega)$ rely on what they observe in the other market.

In this mixture model of the evoked set, the distribution of the option (μ^H, Δ^H) in the evoked set is thus as follows:

- For responders who were responders in the PC market and observed only that market, $(\mu^H, \Delta^H) = (1, 2/3)$ with probability 1.
- For responders who were responders in the RC market and observed only that market, $(\mu^H, \Delta^H) = (0, -1/3)$ with probability 1.
- For responders who were proposers in the PC market, $(\mu^H, \Delta^H) = (1, 2/3)$ with probability q and $(\mu^H, \Delta^H) = (0, -1/3)$ with probability $1 - q$.
- For responders who were proposers in the RC market, $(\mu^H, \Delta^H) = (1, 2/3)$ with probability $1 - q$ and $(\mu^H, \Delta^H) = (0, -1/3)$ with probability q .
- For responders who were in the PC market but observed both markets, $(\mu^H, \Delta^H) = (1, 2/3)$ with probability $1 - q + q\omega$ and $(\mu^H, \Delta^H) = (0, -1/3)$ with probability $q(1 - \omega)$.
- For responders who were in the RC market but observed both markets, $(\mu^H, \Delta^H) = (0, -1/3)$ with probability $1 - q + q\omega$ and $(\mu^H, \Delta^H) = (1, 2/3)$ with probability $q(1 - \omega)$.

We now have the following result:

4. Suppose that $\delta\beta < 11$. Define $M_h = \frac{150\beta}{\delta+2\beta}$ and $M_l = \bar{\delta}^2 M_l$, where $\bar{\delta}^2 := \frac{\delta^2+2\delta\beta}{1+2\delta\beta} < 1$. Then
1. For responders who were responders in the PC market and observed only that market, the average MAO is M_h
 2. For responders who were responders in the RC market and observed only that market, the average MAO is $M_l = \bar{\delta}^2 M_h$
 3. For responders who were proposers in the PC market, the average MAO is $qM_h + (1 - q)M_l = M_h [q + (1 - q)\bar{\delta}^2]$
 4. For responders who were proposers in the RC market, the average MAO is $(1 - q)M_h + qM_l = M_h [1 - q + q\bar{\delta}^2]$
 5. For responders who were in the PC market but observed both markets, the average MAO is $(1 - q + q\omega)M_h + q(1 - \omega)M_l = M_h [1 - q + q\omega + q\bar{\delta}^2 - q\omega\bar{\delta}^2]$
 6. For responders who were in the RC market but observed both markets, the average MAO is $(1 - q + q\omega)M_l + q(1 - \omega)M_h = M_h [q - q\omega + \bar{\delta}^2 - q\bar{\delta}^2 + q\omega\bar{\delta}^2]$.

D Convergence results

In this appendix, we show that if players enter the same environment with different experiences, then preferences will eventually converge. We make several simplifications: we assume that $\alpha = 0$, and we assume that the effect of past experience is solely through the reference point. We consider play in periods $t = -T, \dots, 0, 1, 2, \dots \infty$. In periods $t = -T, \dots, 0$, players participate in some n player game (possibly one of the market games), while in periods $t = 1, 2, \dots$ players participate in a non-competitive ultimatum game. As in section 3.3, we let μ_i^t denote the share of the pie that player i received in period t , and set $\mu_i^t = 0$ if all n players received zero payoffs in the respective period. We assume that the proposer's payoff in periods $t > 1$ is given by $k(Y - a)$, while the responder's payoff is given by a , where a is the offer.

In period $t = -T$, player i 's reference point in an n person game is given by $\mu_i^{-T} = 1/n$. In periods $t > -T$, the reference point of player i in an n -player game is given by

$$r_i^t = (1 - \gamma)(1/2) + \gamma \frac{\sum_{\tau=-T}^{t-1} w_{t-1-\tau} \mu_i^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}}$$

where w_0, w_1, \dots is an infinite sequence given by $w_0 = 1$ and $w_j = \delta^j$ for some $\delta \in [0, 1]$

In the simple specification adopted here, the reference point is a convex combination of the "neutral reference point" $1/n$ and the weighted average of past personal payoff experience. Augmenting the specification to allow for observational experience would not change our results.

We consider the evolution of play between a proposer and a responder in periods $t > 0$. We let r_P^t and r_R^t denote the proposer's and responder's period $t > 0$ reference points. We assume that each period, proposers and responders have perfect information about each others' reference points, and play an SPE of the non-competitive ultimatum game. We let M^t denote the minimal acceptable offer of a responder i in period $t > 0$, and let a^t denote the proposer's period $t > 0$ offer.

Throughout this analysis, we will be concerned with steady state preferences and strategies:

1. A steady state is a pair of strategies (a^*, M^*) and reference points (r_P^*, r_R^*) such that

1. (a^*, M^*) is an SPE of the ultimatum game in which players have the fairness reference points (r_P^*, r_R^*)
2. $r_P^* = (1 - \gamma)(1/2) + \gamma \frac{\pi_P^*}{\pi_P^* + \pi_R^*}$ and $r_R^* = (1 - \gamma)(1/2) + \gamma \frac{\pi_R^*}{\pi_P^* + \pi_R^*}$, where π_P^* and π_R^* are the proposer's and responder's steady state SPE payoffs

Our first result in this section is that there is a unique steady state to which play always converges:

5. *Assume that $\gamma < 1$. Then there is a unique steady state $\langle (a^*, M^*), (r_P^*, r_R^*) \rangle$. In the steady state, $a^* > 0$, $a^* < k(Y - a^*)$, and $a^* = M^*$. Moreover, this steady state is globally stable. That is, for any set of initial experiences $\{\mu_i^t\}_{t=-T}^0$, preferences and strategies converge to the steady state:*

$$\begin{aligned} \lim_{t \rightarrow \infty} r_P^t &= r_P^* & \text{and} & & \lim_{t \rightarrow \infty} r_R^t &= r_R^* \\ \lim_{t \rightarrow \infty} a^t &= a^* & \text{and} & & \lim_{t \rightarrow \infty} M^t &= M^* \end{aligned}$$

Proposition 5 shows that if players have enough experience in the ultimatum game environment, then their fairness preferences in that environment can be characterized as a fixed point of an adjustment dynamic. In fact, Proposition 5 shows that our model uniquely pins down what the steady-state fairness preferences can be—the steady state is unique. The only assumption needed to guarantee uniqueness is that $\gamma < 1$: that is, that players' fairness preferences are not completely (though perhaps arbitrarily close to) determined by past experience.

A second prediction of the model is that when players have extreme past experiences as in our market conditions, convergence to the steady state will be monotonic. That is, PC responders should monotonically decrease their MAO's, while RC responders should monotonically increase their MAO's:

6. *Assume that $\gamma < 1$ and that $\frac{\sum_{t=-T}^0 \mu_R^t}{T+1} + \frac{\sum_{t=-T}^0 \mu_P^t}{T+1} \leq 1$.
If $\frac{\sum_{t=-T}^0 \mu_R^t}{T+1} < r_R^*$, then for all $t > 0$, $M^t < r_R^*$ but is strictly increasing in t .
If $\frac{\sum_{t=-T}^0 \mu_R^t}{T+1} > r_R^*$, then for all $t > 0$, $M^t > r_R^*$ but is strictly decreasing in t .*

Proposition 6 simply says that even though responders' MAO's should not reach steady state levels in a finite number of periods, the effect of past market experience should still diminish over time.

Proof of Proposition 5

Step 1: We first show that there is a unique steady state. In any steady state, we must have

$$M^* - \beta[r_R^*(k(Y - M^*) + M^*) - M^*] = 0, \tag{12}$$

which can be rearranged to show that

$$\frac{M^*}{k(Y - M^*) + M^*} = \frac{\beta r_R^*}{1 + \beta}. \quad (13)$$

Offering $a^* = M^*$ is clearly optimal for the proposer, conditional on making an offer that the responder will accept. Moreover, since $r_P^* + r_R^* = 1$ by definition, some algebra shows that

$$r_P^*[k(Y - M^*) + M^*] < k(Y - M^*),$$

from which it follows that the proposer derives positive utility from making an offer $a^* = M^*$. Thus the proposer's optimal strategy is to offer $a^* = M^*$ in any steady state.

Plugging in $a^* = M^*$ into (13), and using the definition of r_R^* , we now have that

$$r_R^* = (1 - \gamma)(1/2) + \gamma \frac{\beta}{1 + \beta} r_R^*. \quad (14)$$

Equation (14) is a linear equation in r_R^* with a unique solution given by

$$r_R^* = \frac{(1 - \gamma) + \beta(1 - \gamma)}{2 + 2\beta(1 - \gamma)}. \quad (15)$$

Thus there can be at most one steady state. We now show that the unique solution does, indeed, correspond to a steady state. First, examination of equation (15) shows that $r_R^* \in (0, 1)$: since $(1 - \gamma) < 2$, it is clear that the numerator is smaller than the denominator. Next, by definition of M^* , accepting an offer of $a^* = M^*$ is weakly optimal for the responder. And as we have already established, offering $a^* = M^*$ is also optimal for the proposer.

Step 2: We now show that for each $\epsilon > 0$, there exists a $t \geq 1$ such that $r_R^t + r_P^t \leq 1 + \epsilon$. To see this, notice that $\mu_R^t + \mu_P^t \leq 1$ for $t \geq 1$, regardless of the outcome in period t . Thus

$$\begin{aligned} r_R^t + r_P^t &= (1 - \gamma) + \gamma \left(\frac{\sum_{\tau=-T}^{t-1} w_{t-1-\tau} \mu_R^\tau + w_{t-1-\tau} \mu_P^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \right) \\ &\leq (1 - \gamma) + \gamma \left(\frac{\sum_{\tau=-T}^0 w_{t-1-\tau} \mu_R^\tau + w_{t-1-\tau} \mu_P^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} + \frac{\sum_{\tau=1}^{t-1} w_{t-1-\tau}}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \right) \\ &= 1 + \gamma \left(\frac{\sum_{\tau=-T}^0 w_{t-1-\tau} \mu_R^\tau + w_{t-1-\tau} \mu_P^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} - \frac{\sum_{\tau=-T}^0 w_{t-1-\tau}}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \right) \end{aligned}$$

But

$$\frac{\sum_{\tau=-T}^0 w_{t-1-\tau} \mu_R^\tau + w_{t-1-\tau} \mu_P^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \leq \frac{\sum_{\tau=-T}^0 2w_{t-1-\tau}}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}}$$

and

$$\frac{\sum_{\tau=-T}^0 w_{t-1-\tau}}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \rightarrow 0$$

as $t \rightarrow \infty$. Thus for each $\epsilon > 0$, there exists a $t \geq 1$ such that $r_R^t + r_P^t \leq 1 + \epsilon$.

Step 3: We now show that there is some $t^\dagger \geq 1$ such that $a^t = M^t$ for all $t \geq t^\dagger$; that is, for all $t \geq t^\dagger$, the proposer derives positive utility from offering M^t and having that offer accepted.

Set $r_P^t = 1 - r_R^t + \epsilon^t$. As in the proof of Proposition 5, we have that $r_R^t[k(Y - M^t) + M^t] > M^t$. Thus

$$\begin{aligned} r_P^t[k(Y - M^t) + M^t] &= (1 - r_R^t + \epsilon^t)[k(Y - M^t) + M^t] \\ &< [k(Y - M^t) + M^t] - M^t + \epsilon^t[k(Y - M^t) + M^t] \\ &= k(Y - M^t) + \epsilon^t[k(Y - M^t) + M^t]. \end{aligned}$$

This means that the proposer's utility from offering M^t is such that

$$u_P^t \geq k(Y - M^t) - \beta \max(\epsilon^t, 0).$$

Moreover, because $r_R^t \leq (1 - \gamma)/2 + \gamma = (1 + \gamma)/2$, it easily follows that

$$M^t = \frac{k\beta r_R^t Y}{1 + \beta(1 - r_R^t) + k\beta r_R^t}$$

is bounded away from Y (for all possible β) as long as $\gamma < 1$. Thus we have that for all t , there is some $c > 0$ such that $k(Y - M^t) \geq c$. By step 2, there is a t^\dagger such that $\beta\epsilon^t < c$ for all $t \geq t^\dagger$. Thus there is a t^\dagger such that $k(Y - M^t) - \beta \max(\epsilon^t, 0) > 0$ for all $t \geq t^\dagger$.

Step 4: We now strengthen step 2 to show that $|r_P^t + r_R^t - 1| \rightarrow 0$. By step 3, we now have that $\mu_R^t + \mu_P^t = 1$ for all $t \geq t^\dagger$. Thus for $t > t^\dagger$,

$$\begin{aligned} r_R^t + r_P^t &= (1 - \gamma) + \gamma \left(\frac{\sum_{\tau=-T}^{t-1} w_{t-1-\tau} \mu_R^\tau + w_{t-1-\tau} \mu_P^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \right) \\ &= (1 - \gamma) + \gamma \left(\frac{\sum_{\tau=-T}^{t^\dagger-1} w_{t-1-\tau} \mu_R^\tau + w_{t-1-\tau} \mu_P^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} + \frac{\sum_{\tau=t^\dagger}^{t-1} w_{t-1-\tau}}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \right) \\ &= 1 + \gamma \left(\frac{\sum_{\tau=-T}^{t^\dagger-1} w_{t-1-\tau} \mu_R^\tau + w_{t-1-\tau} \mu_P^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} - \frac{\sum_{\tau=-T}^{t^\dagger-1} w_{t-1-\tau}}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \right) \end{aligned}$$

But since

$$\frac{\sum_{\tau=-T}^{t^\dagger-1} w_{t-1-\tau} \mu_R^\tau + w_{t-1-\tau} \mu_P^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \rightarrow 0$$

and

$$\frac{\sum_{\tau=-T}^{t^\dagger-1} w_{t-1-\tau}}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \rightarrow 0$$

as $t \rightarrow \infty$, it follows that $r_R^t + r_P^t \rightarrow 1$ as $t \rightarrow \infty$.

Step 5: We now finish off the proof of the proposition by proving that the steady state identified in Step 1 is globally stable.

Define $\nu_R^t = \frac{\sum_{\tau=-T}^{t-1} w_{t-1-\tau} \mu_i^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}}$. Define the map $\xi : \mathbb{R} \rightarrow \mathbb{R}$ as follows:

$$\xi(\nu) = (1 - \gamma)/2 + \gamma\nu.$$

Define the map $\psi : \mathbb{R} \rightarrow \mathbb{R}$ as follows:

$$\psi(\nu) = \frac{\beta\xi(\nu)}{1 + \beta}.$$

Notice that ψ is linear in ν and has slope $\gamma\beta/(1 + \beta) < 1$; thus ψ is a contraction and has a unique fixed point. In a steady state, $r_R^* = \xi(\nu_R^*)$, and thus equation (13) implies that

$$\frac{M^*}{k(Y - M^*) + M^*} = \psi(\nu_R^*). \quad (16)$$

But since $\nu_R^* = \frac{M^*}{k(Y - M^*) + M^*}$ by definition, it follows that the unique fixed point of ψ corresponds to the unique steady state.

Now for t^\dagger defined as in step 3, $r^t = \xi(\nu^t)$ and $\frac{M^t}{k(Y - M^t) + M^t} = \psi(\nu_R^t)$ for all $t \geq t^\dagger$. Because ξ is strictly increasing, each value of ν_R^t corresponds to a unique value of M^t . Because ψ is strictly increasing and because $\frac{M^t}{k(Y - M^t) + M^t}$ is strictly increasing in r_R^t , each value of ν_R^t also corresponds to a unique value of M^t . Because ξ and ψ are both continuous functions of ν , showing that $\nu_R^t \rightarrow \nu_R^*$ will thus imply that $M^t \rightarrow M^*$ and $r_R^t \rightarrow r_R^*$. Moreover, since $|r_R^t + r_P^t - 1| \rightarrow 0$ by Step 4, convergence of r_R^t will also imply convergence of r_P^t . And finally, since Step 3 shows that $a^t = M^t$ for all $t \geq t^\dagger$, $\nu_R^t \rightarrow \nu_R^*$ will thus also imply that $a^t \rightarrow a^*$.

Because ψ is an increasing and linear function of ν^t that crosses the 45-degree line exactly once, it thus follows that $\psi(\nu) \in (\nu^*, \nu)$ for $\nu > \nu^*$ and $\psi(\nu) \in (\nu, \nu^*)$ for $\nu < \nu^*$. By definition,

$$\nu_R^{t+1} = \frac{w_0}{\sum_{\tau=-T}^t w_{t-\tau}} \mu_R^t + \left(1 - \frac{w_0}{\sum_{\tau=-T}^t w_{t-\tau}}\right) \nu_R^t \quad (17)$$

is a convex combination of ν_R^t and $\psi(\nu_R^t) = \frac{M^t}{k(Y-M^t)+M^t} = \mu_R^t$, which implies that $\nu_R^{t+1} \in (\nu_R^*, \nu_R^t)$ if $\nu_R^t > \nu_R^*$. Similarly, it follows that $\nu_R^{t+1} \in (\nu_R^t, \nu_R^*)$ if $\nu_R^t < \nu_R^*$.

For t^\dagger defined as in step 3, a simple induction thus implies that if $\nu_R^{t^\dagger} < \nu_R^*$, then ν_R^t will be strictly increasing for $t \geq t^\dagger$ and bounded from above by ν^* . Similarly, if $\nu_R^{t^\dagger} > \nu^*$, then ν_R^t will be strictly decreasing for $t \geq t^\dagger$ and bounded from below by ν_R^* . Because any monotonic and bounded sequence converges, ν_R^t must converge to some $\nu^{**} \in [0, 1]$. Because each value of ν_R^t corresponds to a unique value of M^t , and because ψ is continuous in ν , there must, therefore, exist some M^{**} such that $M^t \rightarrow M^{**}$. Thus

$$\lim_{t \rightarrow \infty} \mu_R^t = \lim_{t \rightarrow \infty} \frac{M^t}{k(Y - M^t) + M^t} = \frac{M^{**}}{k(Y - M^{**}) + M^{**}}.$$

It is then easy to show that

$$\nu^t = \frac{\sum_{\tau=-T}^{t-1} w_{t-1-\tau} \mu_i^\tau}{\sum_{\tau=-T}^{t-1} w_{t-1-\tau}} \rightarrow \frac{M^{**}}{k(Y - M^{**}) + M^{**}}.$$

On the other hand,

$$\psi(\nu_R^t) = \frac{M^t}{k(Y - M^t) + M^t} \rightarrow \frac{M^{**}}{k(Y - M^{**}) + M^{**}}.$$

But since ψ is continuous, we therefore have that $\psi(\nu^{**}) = \nu^{**}$. And because ψ has a unique fixed point, it must be that $\nu^{**} = \nu_R^*$, thus completing the proof.

Proof of Proposition 6

Since $r_P^t + r_R^t \leq 1$ for all $t \geq 1$, the reasoning of Step 3 in the proof of Proposition 5 implies that the proposer will offer $a^t = M^t$ in all periods $t \geq 1$. Thus for $t \geq 1$, $r^t = \xi(\nu_R^t)$ and $\frac{M^t}{k(Y-M^t)+M^t} = \psi(\nu_R^t)$.

As in the proof of Proposition 5, a simple induction thus implies that if $\nu_R^1 < \nu_R^*$, then ν_R^t will be strictly increasing for $t \geq 1$ and bounded from above by ν^* . Similarly, if $\nu^1 > \nu_R^*$, then ν_R^t will be strictly decreasing for $t \geq 1$ and bounded from below by ν_R^* . But since M^t is a monotonic function $\zeta(\cdot)$ of ν_R^t such that $M^* = \zeta(\nu_R^*)$, the result follows.

E Comparing Phase 1 and Phase 2 proposer offers

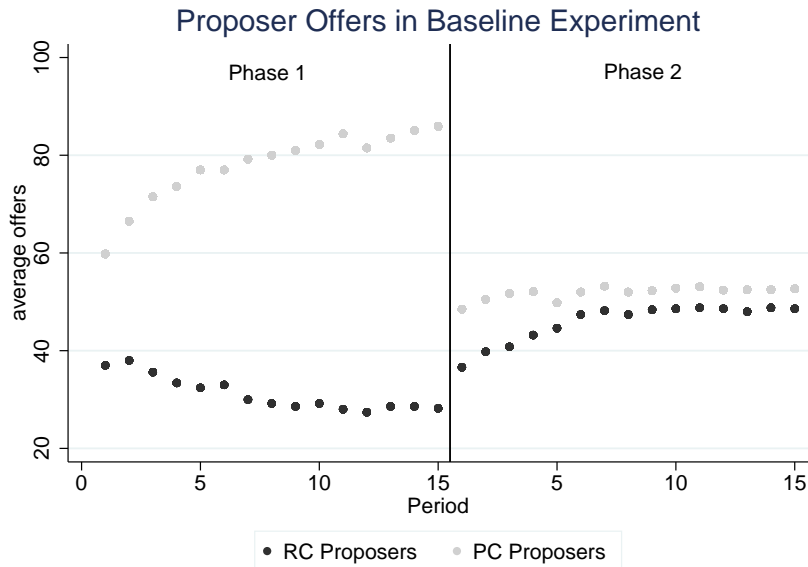


Figure 7: Proposer offers in phase 1 and phase 2 of the Baseline Experiment.

F Regressions with differently coded experiences

In this section, we analyze whether our results on the effect of observational experience and personal payoff experience are sensitive to their respective definitions. First, we have measured observational experience and personal payoff experience simply as the respective average over all 15 periods of phase 1. As an alternative to plain averaging, we consider weighted averages, $G(\boldsymbol{\mu}^\tau, \boldsymbol{\nu}^\tau) = g(\sum \delta^t \mu^t / \sum \delta^t, \sum \delta^t \nu^t / \sum \delta^t)$, which give more weight to more recent periods in phase 1 of the experiment. Second, we have made assumptions about observational experience in the Full Information treatment experiment as well as about personal payoff experience in case of rejection, and we would like to check whether our results are robust to changes in these assumptions.

Table 9 shows results of the same IV GMM regression that is presented in columns (5) and (6) of table 6 of the paper, but uses alternative measures of personal payoff experience and observational experience.

Two different weights have been used to construct the geometric averages: $\delta = 0.9$ and $\delta = 0.95$. It can be seen that using geometric averages does not qualitatively alter our results. Observational experience and personal payoff experience remain significant determinants of responders minimum acceptable offers in phase 2 of the experiment. Also, the magnitude

Table 9: The effect of discounted experiences on minimum acceptable offers

	(1)	(2)	(3)	(4)
	MAO	MAO	MAO	MAO
Observational experience ($\delta = 0.9$)	23.84*** (6.33)	14.73*** (4.62)		
Personal payoff experience ($\delta = 0.9$)	11.76** (5.14)	10.12** (5.02)		
Observational experience ($\delta = 0.95$)			24.85*** (6.67)	15.31*** (4.84)
Personal payoff experience ($\delta = 0.95$)			12.15** (5.32)	10.46** (5.18)
Constant	32.03*** (4.74)	39.44*** (3.15)	31.89*** (4.79)	39.32*** (3.17)
Adj. R^2	0.02	0.05	0.02	0.05
Observations	222	3330	222	3330

Observational experience and personal payoff experience are geometric means of first period offers and payoffs. The discount rate is either $\delta = 0.9$ or $\delta = 0.95$. All regressions are IV GMM regressions, similar to those in table 6. Observational experience and personal payoff experience are instrumented using 6 dummies, one for each market treatment and experiment. Columns (1) and (3) contain period 1 observations of phase 2 only. Columns (2) and (4) use data from all 15 periods. Standard errors are clustered at the phase 1 market matching group level (30 clusters). Significance levels: * = 10%, ** = 5% and *** = 1%.

of the estimated coefficients remains quite stable. Consequently, our results on the impact of observational experience and personal payoff experience are robust to alterations in the construction of these measures.

Second, we had to make assumptions on experienced payoff shares in case of rejections, i.e., when the total sum of payoffs is 0. In the paper, we have assumed that in these cases, the personal payoff experience is equal to 0. An alternative intuition is that in an N -player group, $\mu_i^t = 1/N$ when $\Pi^t = 0$, to reflect the possibility that when everyone gets the same payoff (even when it's zero) the player feels like it was such an equitable outcome that his subsequent feelings of entitlement move towards him getting an even share of the surplus. Additionally, in the Full Information experiment, all subjects received feedback about the average offer in the RC and in the PC market. We have assumed that all subjects correctly weight the information from the PC market twice as much as the information from the RC market, reflecting the fact that there are twice as many offers comprised in the average offer of the PC market. Alternatively, it is possible that subjects weigh these two pieces of information equally. Consequently, we test the robustness of our results with regard to the assumed observational experience of subjects in the Full Information experiment.

Table 10 replicates columns (5) and (6) of table 6 from the main paper using these alter-

Table 10: Alternative codings of observational experience and personal payoff experience

	(1)	(2)	(3)	(4)	(5)	(6)
	MAO	MAO	MAO	MAO	MAO	MAO
Observational experience2	26.00*** (7.06)	15.96*** (5.11)			26.72*** (6.93)	16.43*** (4.97)
Personal payoff experience	12.59** (5.52)	10.84** (5.36)				
Observational experience			26.72*** (6.93)	16.43*** (4.97)		
Personal payoff experience2			12.40** (5.48)	10.92** (5.34)	12.41** (5.48)	10.92** (5.33)
Constant	31.71*** (4.84)	39.16*** (3.19)	31.26*** (4.74)	38.83*** (3.15)	31.26*** (4.74)	38.83*** (3.15)
Adj. R^2	0.025	0.047	0.030	0.050	0.030	0.051
Observations	222	3330	222	3330	222	3330

Observational experience2 differs from observational experience in how observed offers in the Full Information experiment are weighted. Observational experience gives 2/3 weight to the average offer in the PC market and 1/3 weight to the average offer in the RC market, reflecting the fact that there are twice as many proposers in the PC market. Observational experience2 weighs the average offers in the RC market and PC market equally in each period. Personal payoff experience2 codes personal payoff experience as 1/3 in case an offer is rejected by all subjects. Personal payoff experience, on the other hand, codes rejections as a personal payoff experience of 0. All regressions are IV gmm regressions, similar to those in table 9. Observational experience and personal payoff experience are instrumented using 6 dummies, one for each market treatment and experiment. Columns (1), (3) and (5) contain period 1 observations of phase 2 only. Columns (2), (4) and (6) use data from all 15 periods. Standard errors are clustered at the phase 1 market matching group level (30 clusters). Significance levels: * = 10%, ** = 5% and *** = 1%.

native codings of personal payoff experience and observational experience. Personal payoff experience2 is the recoded personal payoff experience variable and observational experience2 is the recoded observational experience variable. Columns (1)-(6) replicate the original estimations in table 6 of the paper using different combinations of the recoded variables. Again, it can be seen that our estimates are robust to these changes in the definition of observational experience and personal payoff experience. The magnitude of the estimated coefficients is remarkably stable, and observational experience and personal payoff experience remain significant determinants of responders' minimum acceptable offers independent of the precise definition of these terms.

G Ruling out Anchoring

As we note in section 4.3 of the paper, experimental evidence has shown that individuals can be influenced by arbitrary anchors (Lichtenstein and Slovic, 2006; Kahneman and Tversky, 2000; Ariely et al., 2003; Simonson and Tversky, 1992), and that behavior that appears to be consistent with expressing a particular preference can in fact be the result of arbitrary anchoring.⁴⁵ On the face of it, our path-dependence account may seem very similar to anchoring. However, there is one crucial difference. We posit that past experience affects *preferences*, and we further show that differences in preferences will affect behavior in environments such as the Ultimatum Game, but that they will not affect behavior in competitive market games as in phase 1 of our experiment (see Appendix A). In contrast, standard anchoring and adjustment theory (Tversky and Kahneman, 1973) does not make such a prediction. This theory states that subjects' choice of action (e.g., offer) starts at some anchor, and then is incompletely adjusted toward the optimal choice of action. Formally, the choice of action is given by $a = (1 - \kappa)\vartheta + \kappa a^*$ where ϑ is the anchor, a^* is the optimal action, and $\kappa \in [0, 1]$ is the degree of adjustment away from the anchor. Such anchoring and adjustment theory would predict that anchors should also have an effect in the market games.

Here, we provide evidence that suggests that behavior in our experiment is not driven by the mere provision of arbitrary anchors. To show this, we exploit a design feature in the Full Information experiment. In this experiment, all subjects received feedback about the average offers in both the PC and the RC markets after every period during phase 1 of the experiment. If responders' acceptance behavior were influenced by the provision of arbitrary anchors, we should observe that responders in the RC market in the Full Information experiment show higher acceptance rates than Responders in the RC market in the Baseline or in the Role Switch experiment.⁴⁶ This is the case because in the Full Information experiment, they are subjected to higher anchors than in the Baseline experiment or in the Role Switch experiment, in which responders only get to observe the offers made by their matched proposer.

Table 11 contains information that tests whether the information provided in the Full information experiment indeed provides an alternative anchor and consequently changes proposers' offers. Moreover, table 12 provides results from a probit regression that test whether responders' acceptance behavior is affected by the altered anchor in the Full Information experiment. It turns out that neither proposers' offers nor responders' acceptance behav-

⁴⁵See, however, Fudenberg et al. (2012) and List et al. (2013) for evidence questioning the robustness of these anchoring effects.

⁴⁶We restrict attention to the RC market because in the PC market, responders almost never reject both offers, and hence there is not enough variance in the data to identify a potential impact.

Table 11: OLS regressions on proposer offers

	(1)	(2)
Full info * PC Market	1.32 (0.98)	1.76 (1.40)
Full info * RC Market	-3.73 (3.61)	2.11 (3.77)
PC Market	41.34*** (3.07)	46.73*** (3.40)
Constant	36.99*** (2.99)	31.15*** (3.17)
Adj. R^2	0.69	0.72
Observations	3330	2430

OLS regression of proposer offers on treatment dummies and interactions. PC Market indicates observations from the Proposer Competition market. Full info*PC Market is an interaction between a Full Information experiment dummy and a proposer competition market dummy. Equivalently, Full info*RC Market is an interaction between a Full Information experiment dummy and a responder competition market dummy. Column (1) contains data from all 3 experiments. Column (2) only compares the Full Information experiment and the Baseline experiment. Standard errors are clustered at the phase 1 market matching group level (30 clusters in columns (1) and (3), 22 clusters in columns (2) and (4)). Significance levels: * = 10%, ** = 5% and *** = 1%.

ior in the Full Information experiment is statistically different from the respective behavior in the Baseline or the Role Switch experiment. Interacting a Full Information experiment dummy with both an RC market dummy and a PC market dummy yields insignificant and economically very small coefficients, implying that proposer and responder behavior in neither the RC nor the PC market was affected by the altered anchor in the Full Information experiment. Consequently, anchoring does not seem to be a driving force of behavior in our experimental setting.

Table 12: Probit regressions on responder acceptance decisions

	(1)	(2)
	Accept	Accept
Maxoffer	0.01*** (0.00)	0.01*** (0.00)
Full info * PC Market	-0.05 (0.05)	-0.01 (0.04)
Full info * RC Market	-0.00 (0.03)	-0.04 (0.04)
PC Market	0.02 (0.04)	-0.04 (0.08)
Adj. R^2	0.23	0.21
Observations	3330	2430

Marginal Effects of a Probit regression of responder acceptance decisions on treatment dummies and interactions. PC Market indicates observations from the Proposer Competition market. full info*PC Market is an interaction between a Full Information experiment dummy and a proposer competition market dummy. Equivalently, full info*RC Market is an interaction between a Full Information experiment dummy and a responder competition market dummy. Maxoffer contains the best offer made to responders in a particular round. In responder competition, there is only one offer. In proposer competition, it is the higher of the two offers made. Column (1) contains data from all 3 experiments. Column (2) only compares the Full Information experiment and the Baseline experiment. Standard errors are clustered at the phase 1 market matching group level (30 clusters in columns (1) and (3), 22 clusters in columns (2) and (4)). Significance levels: * = 10%, ** = 5% and *** = 1%.

H The Full Information Experiment

Table 13: Minimum Acceptable Offers in the Full Information Experiment

	OLS		IV	
	(1) MAO	(2) MAO	(3) MAO	(4) MAO
PC Responder	11.55* (6.11)	5.64 (5.64)		
Personal payoff experience			19.69 ** (9.97)	9.62 (9.27)
Constant	48.45*** (4.67)	47.58*** (3.90)	47.17 *** (4.81)	46.95 *** (3.98)
Adj. R^2	0.01	0.03		
Observations	87	1305	87	1305

Columns (1) and (2) show results of an OLS regression. Regressions include a dummy for proposer phase 1 market experience. Column (1) uses data from period 1 of phase 2 only. Column (2) uses all data. Columns (3) and (4) show IV GMM regressions, in which payoff experience is instrumented using 6 dummies, one for each market treatment and experiment. Observational experience cannot be included here because it is constant across treatments in the Full Information experiment. Column (3) contains data from period 1 only. Column (4) contains data from all 15 periods. Standard errors are clustered at the phase 1 market matching group level (18 clusters). Significance levels: * = 10%, ** = 5% and *** = 1%.

I Phase 1 of the Role Switch and the Full Information Experiments

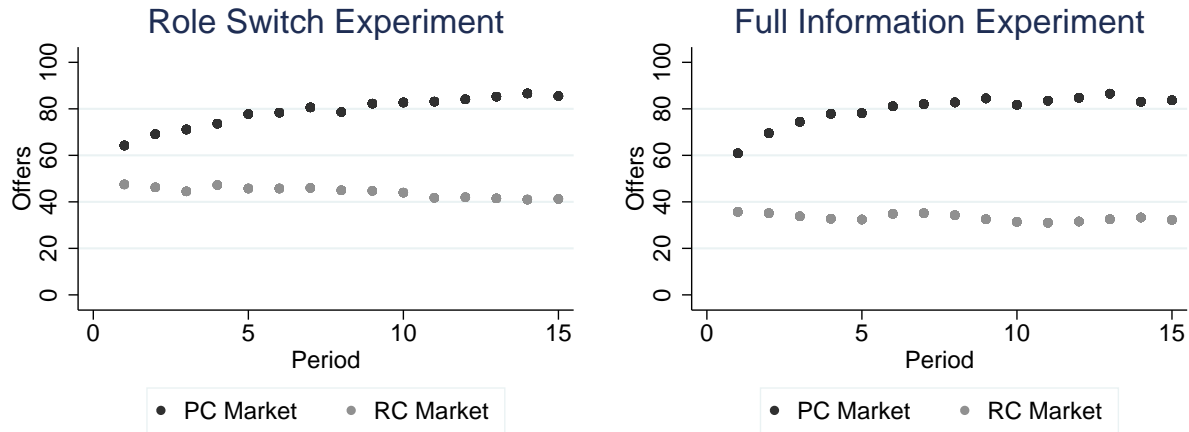


Figure 8: Left Panel: Average offers over time under responder competition and under proposer competition in phase 1 of the Role Switch experiment. Right Panel: Average offers over time under responder competition and under proposer competition in phase 1 of the Full Information experiment

The exogenous assignment to either the PC market or the RC market also had strong effects on market outcomes and personal payoff experiences in the Role Switch and the Full information experiment. Figure 8 shows average offers in phase 1 of both experiments for all 15 periods. Not surprisingly, a very similar pattern to phase 1 of the Baseline experiment emerges. In both experiments, the difference in average offers between the RC market is large. In an OLS regression of offers on a PC market dummy, offers on average differ by 35 chips in phase 1 of the Role Switch experiment and by 46 chips in the Full Information experiment. Both differences are highly significant ($p < 0.001$).