What Makes a Price Fair? An Experimental Study of Market Experience and Endogenous Fairness Norms*

Holger Herz and Dmitry Taubinsky

October 24, 2014

Abstract

Evidence shows that people's fairness preferences are an important constraint for what constitutes an acceptable economic transaction. This paper makes three contributions: First, we provide clean evidence that the transactions people experience in markets in turn shape what is considered acceptable, implying a bidirectional relationship between markets and fairness preferences: Buyers used to high prices are more likely to perceive high prices as fair than buyers used to low prices. Second, we decompose market experience into two different components and separately quantify their impact on this path dependence:An individual's *personal payoff experience* directly captures the personal transaction surplus realized in market transactions. *Market observation* captures all observational information that is contained in market outcomes. We show that both components matter, and independently shape fairness preferences. Third, we provide a simple and tractable model of path-dependent fairness, and estimate its structural parameters. The estimates imply a substantial deviation from classical fairness models that assume exogenous fairness criteria. Our results have implications for price discrimination, labor markets, and dynamic pricing.

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

JEL Classification Codes: C78, C91, D01, D03,

*We thank Philippe Aghion, Bjrn Bartling, Tom Cunningham, Stefano DellaVigna, Ernst Fehr, Oliver Hart, Michael Kosfeld, David Laibson, Michael Luca, Sendhil Mullainathan, Matthew Rabin, Al Roth, Klaus Schmidt, Florian Zimmermann, as well as seminar participants at various universities and conferences 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." Buyers used to high prices are more likely to perceive high prices as fair than buyers used to low prices, and similarly, employees used to low wages are more likely to perceive low wages as fair.[3]

Our paper makes three contributions. First, we provide direct evidence for the hypothesis that past market experience shapes future feelings of entitlement. Second, we decompose market experience and separately quantify two broad mechanisms generating this path dependence: personal payoff experience and market observation. Personal payoff experience relates directly to the individually experienced payoffs resulting from a transaction. Market observation captures the information contained in market transactions, which is independent of an actor's role in the market (and does not even require direct involvement in the market

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

[2]See, 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. Others have argued that social norms may endogenously arise from equilibrium selection in repeated games (Binmore and Samuelson, 1994).

[3]Other determinants of people's feeling of entitlement are exposure to different kinds of bargaining environments (Binmore et al., 1991), or entitlements that stem from investments into the production of the surplus that is to be devided between parties (for reviews of the evidence on such factors, see Handgraaf et al. (2003) or Zwick and Mak (2012)). We differ from these literatures in that we look at history dependence, and not context dependence of fairness preferences.

transaction).[4] Third, our paper combines a simple model with experimental data to structurally estimate the parameters governing preference formation. These structural estimates facilitate a quantitative comparison with fairness models with static reference points such as those advanced by Fehr and Schmidt (1999), Bolton and Ockenfels (2000), Charness and Rabin (2002) and others. Taken together, our results - which involve a total of three experiments and 444 subjects - imply considerable path-dependence in fairness reference points, which constitutes a substantial deviation from these existing theories.

Path dependence in people's fairness preferences has immediate economic implications. First, 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. Similarly, the abolishment of an above market clearing minimum wage would not result in an immediate downward adjustment of wages (Falk et al., 2006). Benjamin (2014) explores the theoretical implications of reference-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 two sources of path dependence reveals new implications for managing price discrimination via information provision.

The broad hypothesis that preferences are malleable in market contexts is supported by recent evidence showing that the mere presence of a market can change people's motives to avoid moral transgressions or to act socially responsible (Falk and Szech, 2013; Bartling et al., 2014), that exposure to violence can increase prosociality toward ingroup members (see, e.g., Voors et al. (2012)), and that contractual agreements between parties serve as a reference for how the completed transaction is evaluated ex-post (Fehr et al., 2011). Aside from the survey evidence in Kahneman et al. (1986), however, there is no direct evidence for the hypothesis that differences in experienced market outcomes will lead to different feelings

[4]In the learning literature, similar differentiations have been made. For example, experience-weighted attraction learning (Camerer and Ho, 1999) assumes that more weight is given to directly experienced outcomes than to counterfactual outcomes when updating attrations of actions in a reinforcement learning type model. Simonsohn et al. (2008) differentiate between "experienced information" and "observed information," and show a differential impact in updating beliefs about opponent play.

[5]See 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.

of entitlement.[6]

Section 2 reports our *Baseline Experiment*, which provides clean evidence for the hypothesis that fairness preferences are shaped by past transactions. In the first phase of our experiment, all subjects participate in one of two market games. In the *proposer competition* (PC) market (Roth et al., 1991), two *proposers* make an offer of a monetary allocation to one responder, who can choose to accept either one or zero of those offers. In the *responder competition* (RC) market (Fischbacher et al., 2009), 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 offers are accepted. Consistent with previous evidence, market conditions have a large impact on the offers: in the PC market, competitive pressures force proposers to give up most of their surplus, while in the RC market, proposers keep most of their surplus.

In the second phase of the experiment, proposers and responders are matched one-on-one in a variant of the ultimatum game (Güth et al., 1982), and proposers again make offers to responders. Consistent with previous studies, responders are 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 are an important reference point for the types of offers they are willing to accept: In period 1 of the ultimatum game, the lowest acceptable offer of a responder who started in the PC market is 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 forces proposers to make very favorable offers to the responders have 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 dissipates by only about one-half of its period-one value.

To further quantify the deviation from static fairness models, we estimate structural fairness parameters governing subjects' behavior in Phase 2 of the baseline experiment. Our estimates imply that responders coming from the RC market feel entitled to an amount that is 20% less than the static fairness norm of equal division assumed by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000); conversely, responders coming from the PC market feel entitled to an amount that is approximately 20% more than the equal division norm.

Having provided clear evidence for path dependence in Section 2, in Section 3 we decompose market experience into two potential sources of path-dependence that may be generating this kind of behavior: *Personal payoff experience* and *market observation.* Market observation corresponds to observational information that is contained in market outcomes. Such information may be an important reference for inferring social norms, which in turn shape

[6] We discuss our contribution and relation to the literature in more detail in section 5.

individual fairness reference points.[7] Personal payoff experience, on the other hand, corresponds to the private surplus that individuals realize in market transactions. Individual's who are used to realize substantial surpluses in market interactions may in turn feel entitled to such surpluses, which implies subsequent reluctance to accept (relatively) unfair treatment. On the other hand, individuals who are used to be deprived in market interactions may become used to such treatment and desensitized to perceived unfairness, which in turn lowers their fairness reference point. In our context, we define market observation as the average offer observed in each round of the market, and personal payoff experience as the average share of surplus a person obtains in each transaction.

In our Baseline experiment, market observation and personal payoff experience are, by design, almost perfectly correlated. Responders in the RC market observe low offers and receive low payments, and vice versa. To separate the potentially differential impact of market observation and personal payoff experience, we report two additional experiments, the *Role Switch experiment* and the *Full Information experiment*. Both experiments are motivated by two economically meaningful ways in which market observation and personal payoff experience differ. First, they may differ substantially for consumers who don't occupy a single role in a market, as when they switch roles between being buyers and sellers. Our *Role Switch experiment* is motivated by this kind of role reversal. This experiment was identical to the Baseline Experiment except with one crucial difference: In Phase 2, subjects who were previously proposers in Phase 1 become responders in Phase 2; whereas subjects who were previously responders in Phase 1 become Proposers in Phase 2.

Second, personal payoff experience and market observation can differ substantially when the price a consumer pays for a particular good is different from the average price posted in the market. This kind of situation is common in markets with publicly observed price discrimination, where all buyers observe the same price distribution, while at the same time experiencing different payoffs. Our *Full Information experiment* is motivated by this divergence between observed and experienced offers. This experiment was identical to the Baseline Experiment except for the following crucial difference: 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.

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 market observation versus personal payoff experience. Interestingly, we find that both matter. We

[7] For example, Bohnet and Zeckhauser (2004) show that informing responders about average offers of all proposers prior to their acceptance decision increases rejections of unfair proposals, consistent with an inferred social norm from market observation. Cohn et al. (2014) show that social comparisions at the workplace affect worker gift-exchange.

confirm this both in our reduced-form estimates, as well as in our structural estimates of a simple model that specifies how personal payoff experience and market observation shape the fairness reference point. Our results show that 50-75% of the path-dependence effect is driven by market observation, and 25-50% is driven by personal payoff experience.

Section 4 examines the robustness of our conclusions to alternative explanations. Our three experiments allow us to rule out a number of potential explanations for our results, including anchoring and classical theories of reference dependent preferences. Section 5 discusses our results in the broader context of the nascent literature on preference formation. Section 6 concludes.

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)[8], and the market game first introduced by Roth et al. (1991). In each 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.[9] 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.

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

[9] We 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.

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[10]. 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 proposer who made the offer and the responder receive chips according to the proposed split. The proposer whose offer was not accepted receives zero 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 were informed about the actual offer, but after the offer was made, responders stated a minimum acceptable offer (MAO) amount; that is, each responder stated 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 was 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 get zero chips if the proposed offer is rejected.[11]

[10] In all games, offers had to be multiples of 5 chips.

[11] Our use of the strategy method in phase 2 but not in phase 1 makes the responders' choice sets very different between the two phases. In phase 1, responders are given choices $A_1 = \{accept, reject\}$, while in phase 2, they are given choices $A_2 = \{0, 5, \ldots, 100\}$. 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 MAOs 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.

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 rematched 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 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).

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, though the pairs were randomly reshuffled every period within each matching group. 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.

Within the chosen phase, 4 periods were selected at random.[12] The points earned in the selected periods were then converted into Swiss Francs, with the exchange rate of points to Swiss Francs set at 10:1.

In total, we ran 5 sessions of the Baseline Experiment, totaling to 150 subjects.[13] 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 Zurich University and the Federal Institute of Technology in Zurich.[14] On average, an experimental session lasted 75 minutes with an average payment of CHF 43.5 ($47.50), including a show-up fee of CHF 15.

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, but we acknowledge that there are a number of nuances that our model does not capture.

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

$$U_i = \pi_i - \beta max\left(r_i\Pi - \pi_i, 0\right) - \alpha \max\left(\pi_i - r_i\Pi, 0\right),$$

where $\Pi = \sum_j \pi_j$ is the total surplus. Here, r is the share of the total surplus a player feels entitled to, or what BO call the "perceived social reference point". Basically, our model is a piecewise-linear version of Bolton and Ockenfels' ERC model. The parameters $\alpha \in [0, 1]$, $\beta \geq 0$ capture, respectively, the disutility associated with player i feeling that he got more or less than what he feels he is entitled to.[15] Both FS and BO focus on equity theories in which all feelings of entitlement derive from a fixed, exogenously given equity norm that is not shaped by past experience. In the context of our specific formal model, this would translate into $r_j = 1/N$ for all players in an N-player game, irrespective of past experiences.

[12] We 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).

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

[14] Subjects were drawn from a database of volunteers using ORSEE (Greiner, 2004).

[15] Following 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.

We, however, do not set r_j equal to $1/N$ but instead test the idea that it may be shaped by past experience.[16]

What are the testable implications of path-dependent fairness reference points? For phase 1 of our experiment, the possibility that $r \neq 1/3$ and is potentially shaped by past experience does not generate sharp testable implications. In the Appendix, 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 perceived entitlements r. In the context of our experimental payoffs, the smallest offer a responder is willing to accept is given by

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

which is a strictly increasing function of r. Thus if experience affects the perceived entitlement r then it should have a direct effect on the minimally acceptable offers (MAOs) of responders.

Our basic hypothesis is that responders who are used to receiving low offers in phase 1 will feel entitled to less than responders who are used to receiving relatively high offers from proposers. As we discuss in more detail in Section 3, where we further specify the adjustment process for the reference points r_j, this effect of experience may come from two related channels: the experience of receiving high or low payoff shares, or the observation of high or low offers.

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

$$3(100-a)Q(a) - \beta max\left[r(300-2a) - 3(100-a), 0\right] - \alpha max\left[r(300-2a) - a, 0\right] \tag{2}$$

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 affects r. Second, all models of fairness assume that people are more concerned

[16] In principle, r_j could be a function of more then just past experience. For example, entitlements have also been shown to be affected by ex-ante investments into the production of the surplus that is to be divided. In principle, our framework could be ammended to also capture such influences on individual fairness reference points. For simplicity and tractability, however, we will solely focus on past experience as a determinant of individual fairness reference points.

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' MAOs. Although our framework is consistent with the idea that phase 1 outcomes may affect proposer behavior in phase 2, studying differences in proposer behavior is a less clean and direct test of the path-dependent fairness hypothesis. We thus focus most of the analysis on responders' MAOs, but come back to exploring proposer behavior in Section 3.5. We will show that our path-dependence hypothesis is consistent both with responder and proposer behavior.

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 7. 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 significant in a regression of offers on a PC market dummy, with standard errors clustered at the phase 1 market level for each session.

Despite these large differences in offers, the right panel of figure 7 shows that the probability that an offer was accepted does not differ that much by treatment. 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 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 plots responders' minimal acceptance thresholds. 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 acceptance thresholds that are 13 chips higher than responders who have previously been in the responder competition market, which translates to PC Responders stating minimum acceptance thresholds that are 36 percent higher than the acceptance thresholds 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.[17] The figure shows that the effect of Phase 1 experience decays only slightly to about 10 chips over the course of the 15 periods.[18] Intuitively, differences in MAOs 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 fairness reference points are expected to converge. We demonstrate this point formally in Appendix C.

2.5 Structural Estimates

The reduced form estimates provide clear evidence of path-dependence of social preferences, but they do not provide clear guidance about the magnitude of path-dependence in terms of structural parameters. 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 the phase 2 entitlements r in equation (1).

Ideally, we would want to estimate the equations

$$\begin{aligned} MAO_{RC,it} &= 100\frac{3\beta r_{RC}}{1+3\beta r_{RC}-\beta r_{RC}+\beta}+\epsilon_{it} \\ MAO_{PC,it} &= 100\frac{3\beta r_{PC}}{1+3\beta r_{PC}-\beta r_{PC}+\beta}+\epsilon_{it} \end{aligned}$$

[17] 15 regressions identical to the regression in column (1) of table 2 were conducted, one for each period. The figure shows the coefficient on the PC Responder dummy.

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

where r_{RC} corresponds to the reference point of RC responders, r_{PC} corresponds to the reference point 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. These equations give us the two moment conditions

$$E\left[MAO_{RC,it} - 100\frac{\beta r_{RC}}{1 + 3\beta r_{RC} - \beta r_{RC} + \beta}\right] = 0 \quad (3)$$

$$E\left[MAO_{PC,it} - 100\frac{\beta r_{PC}}{1 + 3\beta r_{PC} - \beta r_{PC} + \beta}\right] = 0 \quad (4)$$

An econometric challenge is that while equations (3) and (4) contain three parameters, our exogenous variation in phase 1 market experience gives us only two possible moment conditions. We thus consider two approaches to reducing the number of parameters. Our first approach is to assume that $r_{RC} = 1/2 - \rho$ and $r_{PC} = 1/2 + \rho$ in equations (3) and (4) and to estimate ρ and β. Our second approach is to fix β at various values close to the estimate from the first approach, and then to estimate r_{RC} and r_{PC}.

Given the two-parameter vector $\xi = (\beta, \rho)$ or, for fixed β, $\xi = (r_{RC}, r_{PC})$, 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 simply finds 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. Columns (3)-(5) estimate r_{RC} and r_{PC} for $\beta = 0.5$, $\beta = 0.6$, and $\beta = 0.7$, respectively, using period 1 data.

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 (as one would expect, given that fairness preferences are predicted to converge (see Appendix C), but still significant at $p < 0.05$. The direct estimates of r_{RC} and r_{PC} in columns (3)-(5) confirm these previous results. The estimated reference points differ sizably in terms of magnitude, and they are statistically different from each other at $p < 0.01$ for all three values of β. 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.

3 Unpacking the Channels of Path-dependence

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 two broad types of experiences generated by our market games: *market observation* versus *personal payoff experience*. We define market observation as the average offer observed in the 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.

Additionally, preferences may also be shaped by one's own, specific personal payoff experience. Regardless of what is the average market 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 used to obtain 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.

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 $\boldsymbol{\mu}_i^\tau = (\mu_i^1, \mu_i^2, \ldots, \mu_i^\tau)$ denote the period τ history of payoff shares. In the case that $\Pi^t = 0$, we set $\mu_i^t = 0$, to reflect our intuition that receiving a zero payoff should lower one's feelings of entitlement.[19]

Second, we formalizemarket observation. Given offers $a_{1,}^t a_{2,}^t \ldots a_n^t$ 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 $\boldsymbol{\nu}_j^\tau = (\nu_j^1, \ldots, \nu_j^{\tau-1})$ denote the period τ history of average observed offers.

We model a responder's period τ reference point as being given by $r_R = G_R(\boldsymbol{\mu}_i^\tau, \boldsymbol{\nu}_i^\tau)$, where G_R is an increasing function of $\boldsymbol{\mu}_i^\tau$and $\boldsymbol{\nu}_i^\tau$.[20] Similarly, we model a proposer's period τ

[19] 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. In the Appendix, we show that our results are nearly identical under this alternative specification.

[20] For two histories $h = (h_1, \ldots, h_\tau)$ and $h' = (h_1', \ldots, h_\tau')$, we say that $h > h'$ if $h_t \geq h_t'$ for all t and

reference point as being given by $r_P = G_P(\boldsymbol{\mu}_i^\tau, 1-\boldsymbol{\nu}_i^\tau)$, where $1-\boldsymbol{\nu}_i^\tau := (1-\nu_i^{1,} \ldots, 1-\nu_i^\tau)$.

A simple form of G that we will consider for our regression analysis is

$$G(\boldsymbol{\mu}^\tau, \boldsymbol{\nu}^\tau) = g\left(\bar{\mu}^\tau, \bar{\nu}^\tau\right) \tag{5}$$

where $\bar{\mu}^\tau = \sum_{t=1}^{\tau} \mu^t/\tau$ and $\bar{\nu^\tau} = \sum_{t=1}^{\tau} \nu^t/\tau$ are the average personal payoff experience and market observation up to period τ. That is, what matters is the average of past experiences and observations.[21] In our Baseline Experiment, these two components of market experience are almost perfectly correlated. Proposer competition generates high offers, which generates both high market observation and high personal payoff experience for responders. To separate the potentially differential impact of market observation 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 market observation 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 working for low wages who eventually transition to the role of employer, compared to employers in a low wage market who eventually transition to the role of an employee. They will share the same market observation – they both participated in a market with low wages – but their personal payoff experiences will likely differ. Our second experiment is motivated by this kind of role reversal.

Our third experiment is motivated by situations in which personal payoff experience and market observation 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 that this price-discrimination is known to everyone in the market. Here, both the \$10 and \$20 consumers will have the same market observation - they both observe the same market prices - but, by definition, they will have different personal payoff experiences.

$h_t > h'_t$ for at least one t.

[21] In the appendix, we also show that all of our empirical results are robust to considering weighted averages of past experiences and observations: $G(\boldsymbol{\mu}^\tau, \boldsymbol{\nu}^\tau) = g\left(\sum \delta^t \mu^t / \sum \delta^t, \sum \delta^t \nu^t / \sum \delta^t\right)$

3.2 Design of Experiments 2 and 3

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 market observation 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.[22]

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 as well as in the RC market.[23] Consequently, in the Full information experiment, market observation 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.

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

[22]In the Role Switch treatment, the show up fee was 10 CHF and subjects' offer screen featured lists that included all possible offers respectively minimal acceptable offers. This is a difference to the Baseline treatment, 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 H.

[23]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.

Information experiment. Sessions lasted approximately 1-1.25 h and subjects on average earned 41.75 CHF (approx. 46 USD), including a 15 CHF show up fee.

3.2.3 The Six Experimental Conditions

Table 4 summarizes how our three experiments allow us to separately identify the effects of price experience and personal payoff experience. Roughly, our three experiments generate six different cells: (high vs. low personal payoff experience) × (high vs. high & low vs. low market 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". In addition, PC responders in the Baseline experiment as well as PC proposers in the Role Switch experiment are categorized into "high market observation", RC responders in the Baseline experiment as well as RC proposers in the Role Switch experiment are categorized into "low market observation". All subjects from the Full information experiment are categorized as having observed both high & low offers.

3.3 Reduced-form Estimates

We begin by summarizing the behavior conditional on the treatment.[24] Figures 4 and 5 show the average minimum acceptance thresholds for the different combinations of market observation and personal payoff experience. Figure 4 splits the data by personal payoff experience. The left panel shows the minimal acceptance thresholds for responders in phase 2 who experienced low payments, conditional on their market observation. Similarly, the middle panel shows minimal acceptance thresholds of responders in phase 2 who experienced high payments, conditional on theirmarket observation. The right panel shows the average effect of market observation, giving equal weight to average observations with low personal payoff experience and with high personal payoff experience.

Figure 5, on the other hand, splits the data by market observation. The three panels from the left show the minimal acceptance thresholds for responders in phase 2 who observed only low, low and high, or only high offers, conditional on the personal payoff experience.

[24] In the Appendix, we show that the Phase 1 outcomes in these two additional experiments are very similar to the Phase 1 outcomes for 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.

The right panel shows the average effect of personal payoff experience, giving equal weight to average observations with low, low and high, and high offer observation.

To statistically analyze the effect of personal payoff experience and market observation separately, we first analyze the Baseline and the Role Switch experiment in isolation, and then analyze the Baseline and Full Information experiments in isolation. The combination of Baseline and Role Switch allows us to focus on the way in which occupying multiple roles in a market can lead to different personal payoff experiences and different market observations. The combination of Baseline and Full Information allows us to focus on the way in which personal payoff experience and market observation may differ when the terms of a person's specific transaction are not representative of the average market outcome.

To jointly estimate the impact of personal payoff experience and market observation on responder MAOs, 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 exogenous instruments for responder personal payoff experience and market observation. 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.[25]

Matching groups with responders who are particularly prone to rejection will drive up offers in Phase 1, which would then create a spurious correlation between Phase 1 offers and Phase 2 MAOs.

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 $\bar{\nu} = \bar{\mu} = \sum_{t=1}^{15} \mu^t/15$ are the average of Phase 1 market observation and personal payoff experience, as in equation (5).

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

[25]To put in another way: for the same reason that we couldn't test our path-dependence hypothesis by simply having subjects play 15 rounds of an ultimatum game and then regressing their MAOs on past experience, we can't simply regress MAOs on past experience even when we have an additional source 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.

Table 6 reports 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 market observation on responders' MAOs. What is also noteworthy is that the effects of past experience and observation do not appear to differ across experiments 2 and 3. The impact of market observation in the Role Switch experiment appears to be very similar in magnitude to the effect of market observation in the Full Information experiment.

Table 7 pools all three experiments for greater power and reports reduced-form estimates of the impact of phase 1 market observation and personal payoff experience on responder MAO's. Column (1) again focusses on period 1, while column (2) uses data from all 15 periods. In these pooled regressions, we again find that both market observation 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 market observation. Moreover, the estimates of the impact of market observation and personal payoff experience are again very similar to the estimates in table 6.

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 acceptance thresholds by approximately 1 chip. Similarly, a 10 percentage points increase in observed average offers increases first period acceptance thresholds 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 coefficient estimates for market observation 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 market observation, individually estimated for every period, similar to the regression for period 1 in column (1) of table 7. Dashed lines indicate 95% confidence intervals of the parameter estimates. Market observation has a large and significant impact on minimal acceptance thresholds 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 Structural Estimates

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 market observation play in shaping the reference point in equation (1). For our structural estimation, we further simplify g by assuming that the reference point is formed through a convex combination of the average market observation and the average personal payoff experience. Thus for responders,

$$g_R(\boldsymbol{\mu}^\tau, \boldsymbol{\nu}^\tau) = (1 - \gamma_{PE} - \gamma_{MO})(1/2) + \gamma_{PE} \sum_{t=1}^{\tau} \mu^t/\tau + \gamma_{MO} \sum_{t=1}^{\tau} \nu^t/\tau \tag{6}$$

where γ_{PE} and γ_{MO} are the weights on personal payoff experience and market observation. Similarly, for proposers,

$$g_P(\boldsymbol{\mu}^\tau, \boldsymbol{\nu}^\tau) = (1 - \gamma_{PE} - \gamma_{MO})(1/2) + \gamma_{PE} \sum_{t=1}^{\tau} \mu^t/\tau + \gamma_{MO} \sum_{t=1}^{\tau} (1 - \nu^t)/\tau \tag{7}$$

We compute a responder's predicted entitlement r as in equation (6), and we use equation (1) to compute a responder's predicted MAO. The six different treatment cells (see table 4) allow us to use Generalized Method of Moments to recover γ_{PE} and γ_{MO} 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 $\xi = (\beta, \gamma_{PE}, \gamma_{MO})$ 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 treatment-session 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 structural model states that the reference point 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 (6). This is somewhat problematic because in phase 2, personal payoff experience and market observation are highly collinear, as they are in experiment 1.[26] 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 market observation is still phase 1 experience. Because of the high degree of phase 2 collinearity, we thus use caution in interpreting the estimates that arise from using all 15 periods of Phase 2 data.

Table 8 presents the results, with column (1) focusing on period 1 only, and column (2) using data from all 15 periods. For the reasons mentioned above, the period 1 only data is a cleaner source of exogenous variation for estimating our three structural parameters.Column (1) shows that both personal payoff experience and market observation receive positive weight in shaping the reference point. While both weights are significantly different from 0, the weight given to market observation, 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} loosing significance, potentially because of the partial collinearity problem.

3.5 Proposer Offers

Last, we turn to Proposer behavior, which should be influenced both by their own fairness motives, and 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 market observation 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 market observation of all responders in the matching group. Table 9 displays regressions analyzing how proposers' and responders' phase 1 experience and observation influence proposers' phase 2 offers. We use all 9 experimental conditions as instruments for the 4 experience and observation variables, and use the iterative GMM estimator as before.

Consistent with the path dependence fairness hypothesis, column (1) shows that in period 1, proposers who are used to receiving a high 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, proposers who are used to observing higher offers, are more likely to make a higher offer. Column (2) of the table uses data from all 15 periods of Phase 2, and shows that proposers'

[26] A 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.

offers are very sensitive to responders' behavior, in the predicted direction: the higher the payoff experience or observed offers of the responders, the higher their MAOs, and thus the higher the offers made by the proposers to these responders. In fact, the responder experience and observation coefficients in column (2) of Table 9 are almost identical to the responder experience and observation coefficients in column (2) of Table 7, suggesting that proposers' offers respond almost one-for-one to responders' MAOs.

In addition to learning over time, Proposers also know to adjust their offers even in period 1 of Phase 2. In our Baseline treatment, while proposers from the PC market offered 84.9 chips in period 15 of Phase 1, they offered only 49.1 chips in period 1 of Phase 2. Similarly, while proposers in the PC market offered 33.1 chips in period 15 of Phase 1, they offered 43 chips in period 1 of Phase 2. This result of an immediate change in offers is consistent with Grosskopf (2003).

4 Related Approaches and Alternative Explanations

We have argued that individual's fairness preferences are path-dependent, and that the fairness reference point is shaped by previous market observation and personal payoff experience. In this section, we discuss whether our experimental findings could be explained by other forms of path dependence that are not grounded in changes in fairness preferences. We discuss other forms of reference-dependent preferences and anchoring. In Section 4.3 we also contribute to the recent debate about whether reference points are determined by the past, such as the status quo, or whether reference points are grounded in expectations.

4.1 Ruling out 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 a reference point, and deviations of outcomes from this reference point are then associated with psychological gain/loss utility, where losses are assumed to have larger psychological consequences than equally sized gains. 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 Ruling out 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 that behavior that appears to be consistent with expressing a particular preference can in fact be the result of arbitrary anchoring. 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, 1974) does not make such a prediction. This theory states that 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. For 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. Such anchoring and adjustment could, in principle, also have an effect in our experiment.

The first thing to note is that only the impact of market observation can possibly be explained by anchoring, similar to the findings by Ariely et al. (2003). Market observation captures all observational attributes of phase 1, such has high and low offers, that could serve as arbitrary anchors. The significant effect of personal payoff experience cannot be explained by such anchoring.

Furthermore, in the appendix, we provide evidence that suggests that behavior in our experiment is generally not driven by the mere provision of arbitrary anchors, i. e., market observation is also unlikely to be the mere consequence of arbitrary anchoring. To show this, we exploit a design feature in the Full Information treatment. Here, 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 an anchor that is substantially lower than offers in the PC market, whereas subjects in the RC market observe an anchor that is substantially higher than offers in the RC markets. As we have already demonstrated, such market observation has large effects on behavior in Phase 2. If subjects' behavior were indeed solely driven by the provision of such anchors, we should observe already in phase 1 that responders in the RC market in the Full Information treatment show higher acceptance rates than responders in the RC market in the Baseline or in the Role Switch treatment.[27] Similarly, we should observe that PC Proposers in the Full Information experiment offer less than PC Proposers in the other two experiments, and that RC proposers in the Full Information experiment offer more than RC proposers in the other two experiments. It turns out that neither acceptance behavior of responders nor offer behavior of proposers in the Full Information treatment is statistically different from acceptance behavior and offer behavior in the Baseline and the Role Switch treatment.

4.3 Adaptive Expectations

In the development of our hypotheses, we have assumed that fairness reference points correspond to past experiences. Alternatively, following Kőszegi and Rabin (2006), it is possible that reference points correspond to expectations. Applied to our setting, a responder may choose to reject a proposer's offer when that offer falls far short of what the responder expected.[28]

A model with *rational* expectations as reference points would not predict our 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.[29]

Another possibility is that a model of *naive*, rather than rational, expectations might be generating the differences in responders' phase 2 behavior. Note, however, that adaptive expectations should be formed solely based on available information about other players' behavior, which is captured by our market observation variable. Personal payoff experience should not play a role in forming naive expectations. This is inconsistent with the reduced form estimates in table 7 and the structural estimates in table 8 which reveal that while the

[27] We 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.

[28] In 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 and Fuster (2011) provide evidence that expectations shape reference points.

[29] Rational 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*.

impact ofmarket observation is larger in magnitude, personal payoff experience also has a significant impact. To further reinforce the point that adaptive expectations can not account for the whole treatment effect in the Baseline experiment, we construct Table 14 in Appendix G. The table analyzes only the Full Information experiment, where all players receive the same information about Phase 1 behavior. We show that the payoff differences generated by Phase 1 market assignment still have a significant impact on behavior in this experiment.

However, 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 market observation.[30][31]We leave it to future work to explore to what extent some of our market observation effect could be determined by naive, backward-looking expectations.

5 Relation to the Literature

Our paper contributes to a nascent literature on preference formation and behavioral spillovers, which we discuss in more detail in this section.

First, our paper contributes to the recent literature on the malleability and stability of preferences. Voors et al. (2012) show that individuals exposed to violence in the civil war in Burundi display more altruistic behavior towards their neighbors. Similar results in other setups also involving violent conflict have been shown by Bauer et al. (2014), Cassar et al. (2012) and Gilligan et al. (2014). Bettinger and Slonim (2006) show that even school intervention programs can affect pupils' altruism towards charities, while Rao (2013) shows that class composition can affect pupils' altruism toward peers of different backgrounds.[32] Our investigation differs from this literature in several ways. First, we study how *outcomes in common economic transactions* impact preferences, rather than rare, life-changing events. In that sense, our results have more immediate implications for the economic analysis of mar-

[30]See Cooper and Dutcher (2011) for a sketch of a model along these lines. As we show in appendix B, however, belief based reciprocity models (Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006) do not naturally give rise to this effect. Evidence that expectations may interact with fairness in other contexts is provided by Gilchrist, Luca, and Malhotra (2014), who find that unexpected bonuses promote more worker reciprocity than expected wage increases.

[31]Alternatively, expectations may not be adaptive but, as in Brunnermeier and Parker (2005), formed endogenously to create higher anticipatory utility. By itself, the Brunnermeier and Parker framework makes no predictions about players' fairness preferences, since people derive utility only from consumption and anticipation of future consumption in that framework. Moreover, even if the framework is combined with a different specification of belief utility, it still cannot predict our results. The simple reason is that because responders' choices are non-strategic, their optimal choice of beliefs will always be a "corner" solution, irrespective of their market experience - if lower expectations increase utility, then responders will choose to expect nothing; if higher expectations maximize utility, then responders will choose to expect everything.

[32]Other papers have analyzed the impact of such shocks on risk and time preferences. See, for example, Callen et al. (2014); Cameron and Shah (forthcoming).

kets. Second, we go beyond a demonstration that preferences are endogenous and distinguish between two mechanisms that might be causing this endogeneity.

Second, our work contributes to a recent literature on the role of ex-ante agreements on the evaluation of ex-post outcomes in bilateral interactions (Hart and Moore, 2008)[33], 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. The reference point in this literature is therefore shaped by prior interaction and contractual agreement between the involved parties. Our notion of reference-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. In fact, Hart and Moore (2008) discuss extensions of their model in which reference points 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.

Third, researchers have been interested in behavioral spillovers across economic games, in particular learning transfers across strategic interactions (Grimm and Mengel, 2012; Bednar et al., 2012; Cason et al., 2011). This literature has shown how beliefs about opponents' play can be influenced by observations of play in similar games. This demonstrates the importance of learning spillovers in coordination games, possibly through belief- or best-response bundling. But different from our results, these papers are perfectly consistent with strategic, profit-maximization behavior and (naive) belief updating. By focusing on the non-strategic decision of responders, we demonstrate spillovers that operate at the level of preferences.

Similarly, 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, Lévy-Garboua et al. (2009) show that history can serve as focal point in multiple equilibria games, and Falk et al. (2006) experimentally show that experimenter-imposed minimum wage laws can cause spillover effects, raising wages even after the removal of the minimum wage law. Finally, Peysakhovich and Rand (2013) 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. Like our work, they are thus able to find a spillover

[33]see also Hart (2009); Fehr et al. (2011, forthcoming); Bartling and Schmidt (forthcoming); Brandts et al. (2012)

effect in a non-strategic interaction. The experimental evidence in these papers is *consistent* with our proposed mechanism of past-dependence fairness preferences.[34] However, we go beyond these papers by not simply demonstrating an effect, but by reporting theoretically motivated experiments that show that the spillover effect is significant and long-lasting even in non-strategic interactions, where it cannot be attributed to belief formation about what strategies maximize the individuals' otherwise exogenous, history-invariant preferences.

6 Conclusion

While most work on social preferences has progressed under the presumption of static preferences, we show that fairness preferences are malleable and endogenously shaped by economic forces through the experience of different market outcomes. We also show that such malleability can be captured in a simple and tractable model with estimable parameters. Our reduced-form and structural results imply significant deviations from existing models of fairness.

Our results highlight the importance of considering reference-dependence in fairness preferences. Existing theories of backwards-looking reference points in the consumption dimension only are sharply inconsistent with our experimental results. And unlike theories of reference dependence that do not invoke fairness preferences - but in line with the insights of Kahneman et al. (1986) - our framework implies that consumers will react very differently to prices increases that are exploitations of market power, as opposed to price increases necessitated by rising costs of production.

In addition to the labor market implications explored by Benjamin (2014) 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 whose fairness reference point is violated, and the long run profits generated through an increased willingness to pay of customers once the reference point has adjusted. This leads to new considerations for dynamic price setting, and may lead to a violation of the Coase Conjecture for intertemporal price discrimination.

Our results on price experience also generate practical implications that could be studied in field settings. For example, a firm trying to price discriminate among consumers should

[34] In all of these papers, a version of our model in which the reference point r is fixed and not shaped by experience would produce multiple equilibria. In Hargreaves-Heap and Varoufakis (2002), for example, both the cooperative equilibrium and the haw-dove equilibrium can be sustained by a standard model of distributional preferences. Similarly, in Falk et al. (2006), because there are several workers making accept-reject decisions, there can be multiple equilibria: If a worker thinks others are likely to reject, then he will reject since then he actually has the ability to punish; if a worker thinks other are not likely to reject then he is better off accepting.

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 model we have introduced could be used to analyze when shrouding price discrimination is payoff maximizing.

More generally, a key implication of our results on price experience is that informing buyers about other prices or workers about other wages should change the kinds of prices or wages the buyers are willing to accept, even when such information is payoff irrelevant. Our results imply that increased information dispersion should have the effect of homogenizing fairness norms. These and other theoretical and empirical extensions of our analysis are directions for future research.

Tables and Figures

	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

Table 1: Overview of Matching Groups

Table 2: The Impact of Responder Experience in the Baseline Experiment

	(1)	(2)
PC Responder	13.000**	10.187*
	(4.208)	(4.921)
PC Proposer		–3.227
		(7.490)
Constant	36.000***	37.413***
	(2.899)	(4.675)
Adj. R^2	0.064	0.041
Observations	75	1125

The regression in column (1) includes minimal acceptance thresholds 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. Robust standard errors are clustered by phase 1 market treatment groups (2 clusters per session, 10 clusters in total). Significance levels: *** = 1%, ** = 5% and * = 10%.

Table 3: Structural Estimates from the Baseline Experiment

	(1)	(2)	(3)	(4)	(5)
ρ	0.106***	0.073**			
	(0.032)	(0.030)			
β	0.668***	0.579***			
	(0.073)	(0.061)			
r_{RC}			0.474***	0.421***	0.383***
			(0.047)	(0.042)	(0.038)
r_{PC}			0.728***	0.647***	0.589***
			(0.063)	(0.056)	(0.051)
Observations	75	1125	75	75	75

Column (1) contains estimates for ρ and β using period 1 data only. Column (2) contains estimates forρ and β using data from all periods. Column (3)-(5) all use data from period 1 only and estimate r_{RC} and r_{PC} using different exogenously given values of β. In column (3), $\beta = 0.5$. In column (4), $\beta = 0.6$. In column (5), $\beta = 0.7$. Robust standard errors are clustered by phase 1 market treatment groups (2 clusters per session, 10 clusters in total). Significance levels: *** = 1%, ** = 5% and * = 10%.

Table 4: Overview of Treatment Variation

	Market Observation		
	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 5: Summary statistics of market observation and personal payoff experience by treatment and experiment

	personal payoff experience	market observation
Baseline PC Responders	0.625	0.568
Baseline RC Responders	0.067	0.140
Role Switch PC Proposers	0.189	0.576
Role Switch RC Proposers	0.700	0.223
Full Information PC Responders	0.660	0.433
Full Information RC Responders	0.069	0.433

Table 6: Comparisons between Experiments 1 & 2 and 1 & 3

	Exp.1 & 2		Exp.1 & 3	
	(1)	(2)	(3)	(4)
market observation	26.685***	15.747***	20.939	21.435*
	(6.999)	(5.313)	(14.802)	(11.601)
personal payoff experience	9.399	11.052	15.565	6.203
	(7.158)	(6.823)	(10.187)	(9.027)
Constant	30.728***	38.755***	35.930***	36.077***
	(5.501)	(3.974)	(3.867)	(3.506)
Adj. R^2	0.034	0.057	0.038	0.038
Observations	135	2025	162	2430

Instrumental variables regressions estimating the impact of Phase 1 personal payoff experience and market observation on Phase 2 behavior. Estimates computed using the iterative GMM estimator. In each column, personal payoff experience and market observation 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. Robust standard errors are clustered by phase 1 market treatment group (2 clusters per session, 18 clusters in total in columns (1) and (2), 22 clusters in total in columns (3) and (4)). Significance levels: *** = 1%, ** = 5% and * = 10%. Columns (1) and (2) additionally contain dummies for high proposer personal payoff experience and high proposermarket observation. Columns (3) and (4) additionally contain a dummy for high proposer personal payoff experience. (including proposer market observation dummies would lead to collinearity, because high and low market observation is unique to the Full Information treatment).

Table 7: Pooled instrumental variables regressions separately estimating the effect of personal payoff experience and market observation on minimum acceptance thresholds in phase 2

market observation	23.928***	15.954***
	(6.568)	(5.105)
personal payoff experience	9.620**	10.849**
	(4.328)	(5.358)
Constant	32.913***	39.162***
	(3.053)	(3.191)
Adj. R^2	0.030	0.046
Observations	222	3330

Instrumental variables regressions estimating the impact of Phase 1 personal payoff experience and market observation on Phase 2 behavior, using data from all three experiments. Estimates computed using the iterative GMM estimator. Payoff experience and market observation are instrumented using 6 dummies, one for each market treatment in each experiment. Column (1) contains period 1 observations of phase 2 only. Column (2) uses data from all 15 periods. All regressions additionally contain dummy variables for high proposer personal payoff experience and for proposer market observation. Robust standard errors are clustered by treatment group (2 clusters pre session, 30 clusters in total). Significance levels: *** = 1%, ** = 5% and * = 10%.

Table 8: Structural Estimates of how market observation and personal payoff experience shape the reference point

	(1)	(2)
γ_{PE}	0.129**	0.059
	(0.056)	(0.071)
γ_{MO}	0.325***	0.269***
	(0.079)	(0.075)
β	0.981***	0.695***
	(0.087)	(0.060)
Hansen's J	4.96	6.16
	(p=0.66)	(p=0.52)
Observations	222	3330

Generalized Method of Moments estimates of the structural model in equation (6). Column (1) restricts to period 1 only, while column (2) uses data from all 15 periods. The six moment conditions are determined by instruments corresponding to the six Experiment × Treatment conditions. Standard errors are computed at the Session × Treatment 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%.

Table 9: Proposer Offers and Experience

	(1)	(2)
market observation	8.584*	7.582
	(4.499)	(4.856)
personal payoff experience	–13.177***	5.879
	(2.585)	(3.687)
avg. responder price experience	5.941	15.975***
	(4.616)	(3.796)
avg. responder personal payoff experience	–5.135	8.563**
	(4.431)	(3.403)
Constant	46.362***	39.362***
	(2.775)	(3.158)
Adj. R^2	0.019	0.105
Observations	222	3330

Instrumental variables regressions of offers on proposer and responder payoff experience and market observation, using data from all three experiments. Both proposer and responder payoff experience and market observation are instrumented using 9 dummies, one for each phase 2 matching group in each experiment. All estimatescomputed 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 treatment group (2 clusters per session, 30 clusters in total). Significance levels: *** = 1%, ** = 5% and * = 10%.

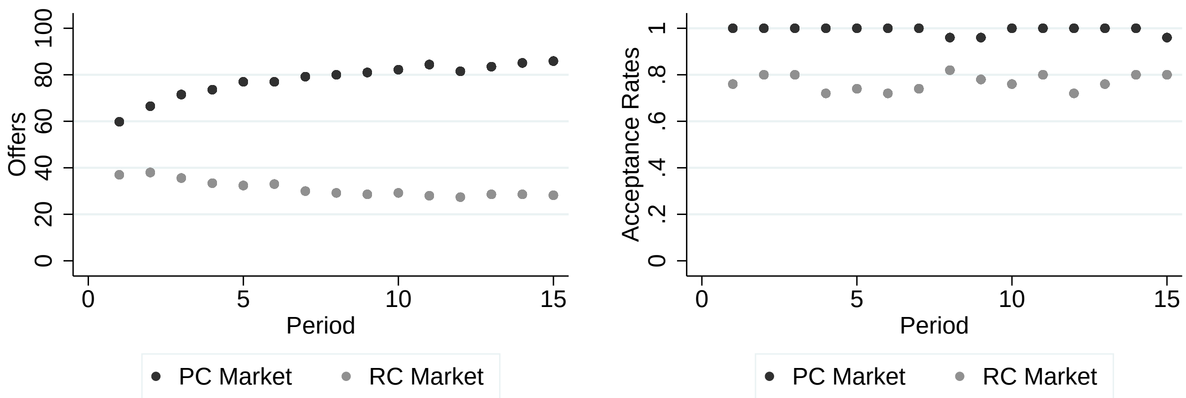

Figure 1: Left Panel: Average offers over time under responder competition (RC Market) and under proposer competition (PC Market) in phase 1 of the Baseline experiment. Right Panel: Acceptance rates of responders over time under responder competition and under proposer competition in phase 1 of the Baseline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

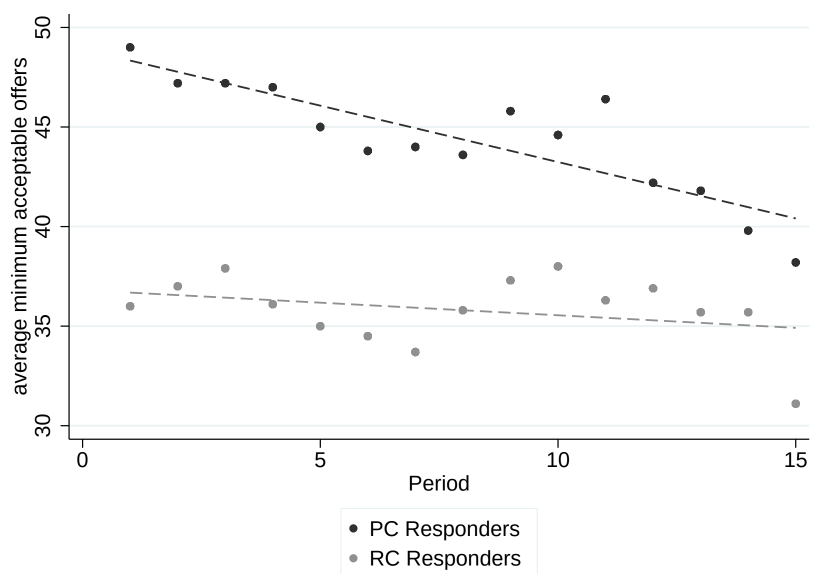

Figure 2: Minimal Acceptance Thresholds of responders. “PC Responders” denotes responders who have participated in the PC market in phase 1. “RC Responders” denotes responders who have participated in the RC market in phase 1. The figure shows average minimum acceptable offers for PC Responders and RC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

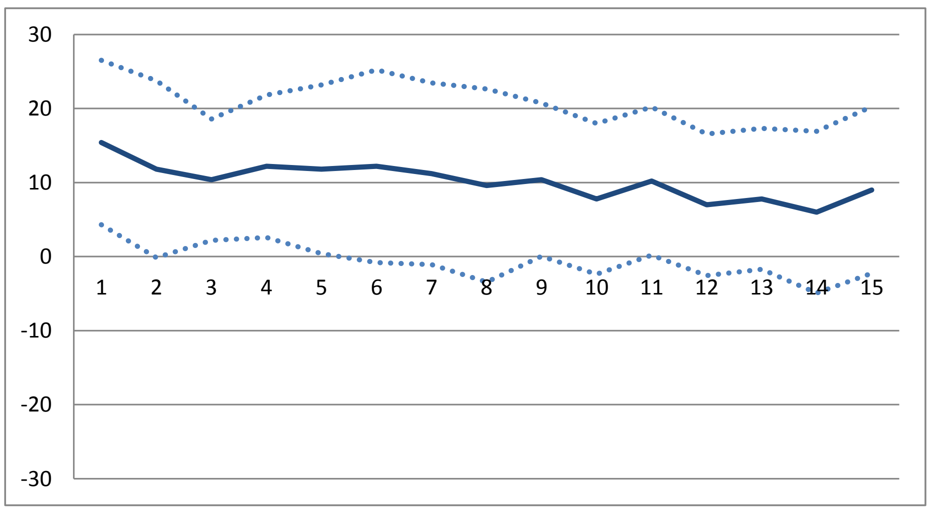

Figure 3: This figure plots the impact that being in the PC market in phase 1 has on responders' MAOs in phase 2, for each of the 15 periods. The regression specification also proposer experience. The dotted lines indicate 95% confidence intervals. Robust standard errors are clustered by the phase 1 market treatment group (2 clusters per session, 10 clusters in total).

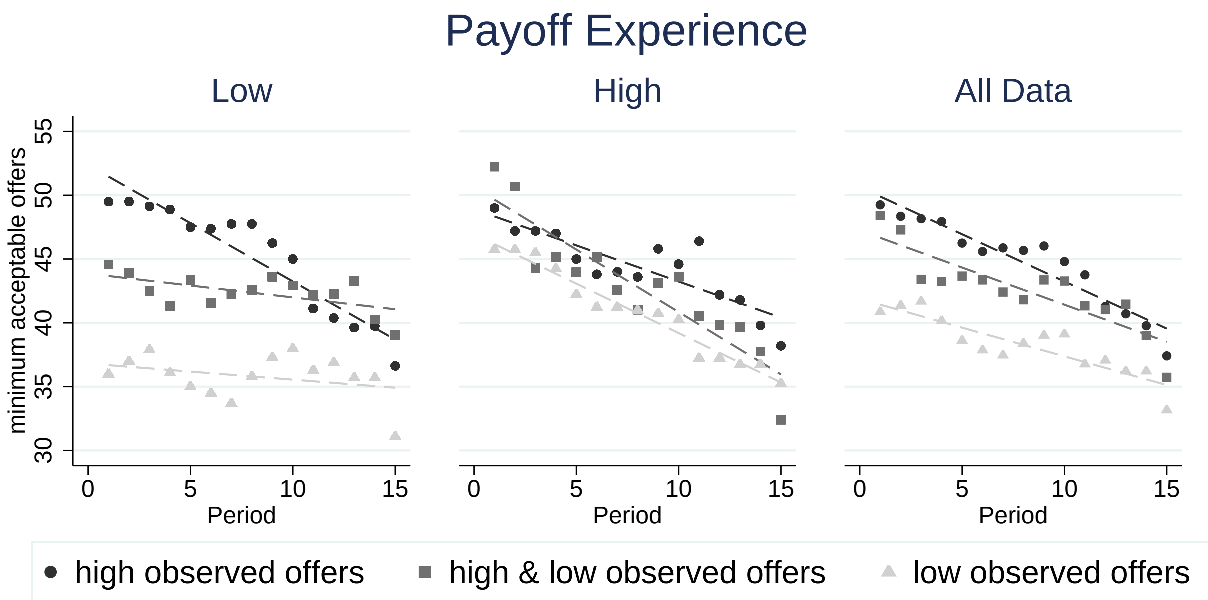

Figure 4: Minimal Acceptable offers by personal payoff experience. All data pools data from all experience treatments, giving equal weight to each experience treatment.

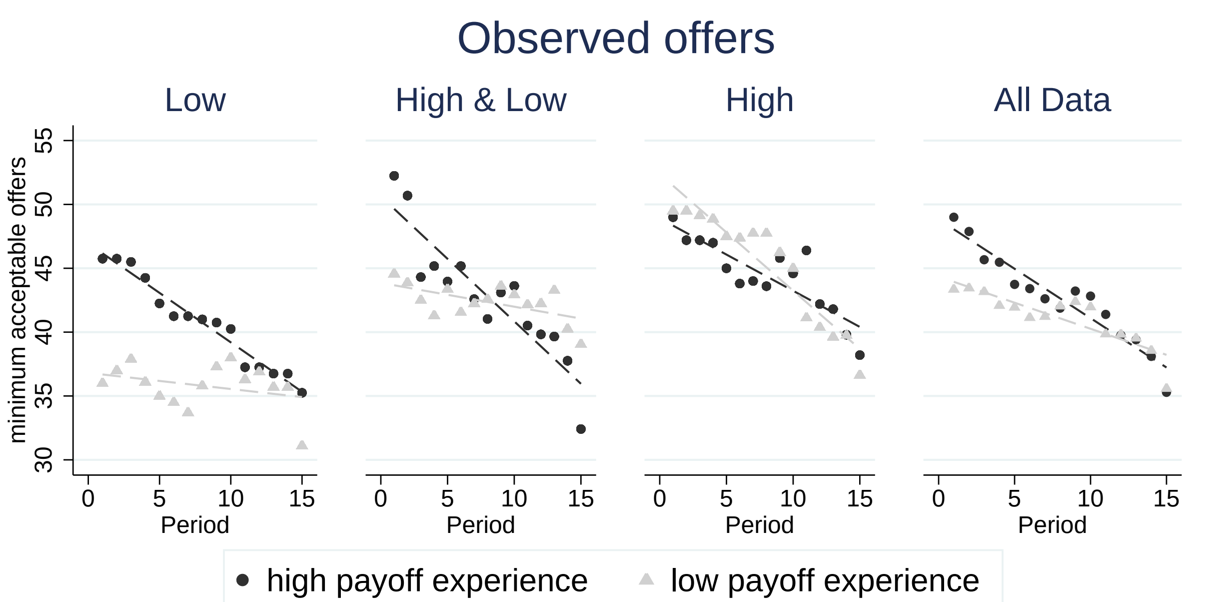

Figure 5: Minimal Acceptable offers by market observation. All data pools data from all observation treatments, giving equal weight to each market observation treatment.

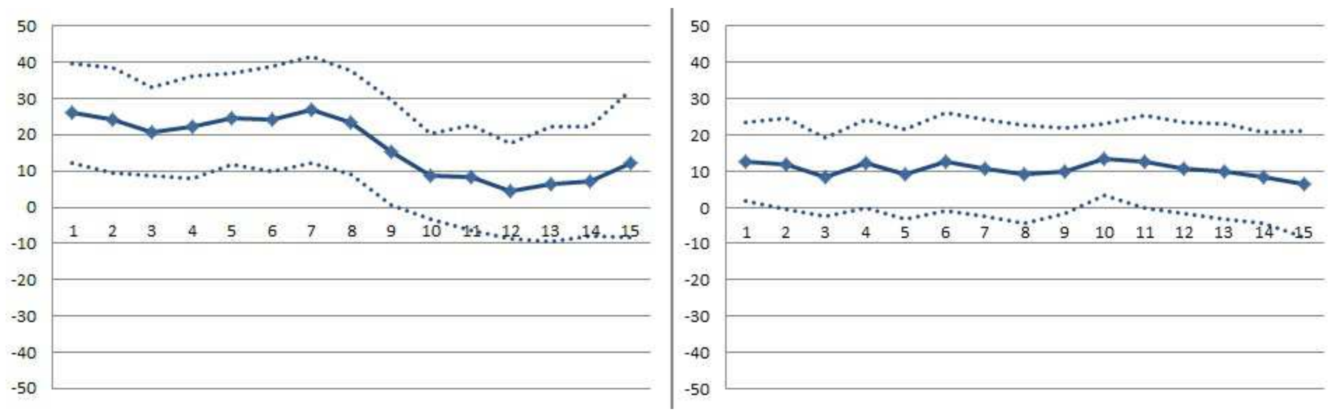

Figure 6: Left Panel: Development of the market observation parameter over time, individually estimated for every period using the same IV regression specification as in column (1) of table 7. Right panel: Development of the personal payoff experience parameter over time, individually estimated for every period. Robust standard errors are clusteredbyt the phase 1 market treatment group (2 clusters per session, 30 clusters in total). The dotted lines indicate 95% confidence intervals.

References

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," ECON - Working Papers 134, Department of Economics - University of Zurich 2014.

Bartling, Björn and Klaus M. Schmidt, "Reference Points, Social Norms, and Fairness in Contract Renegotiations," *Journal of the European Economic Association*, forthcoming.

Bauer, Michal, Alessandra Cassar, Julie Chytilová, and Joseph Henrich, "Wars 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.

Benjamin, Daniel J., "A Theory of Fairness in Labor Markets," *Working paper*, 2014.

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.

Binmore, Ken and Larry Samuelson, "An economist's perspective on the evolution of norms," *Journal of Institutional and Theoretical Economics*, 1994, pp. 45–63.

_ , Peter Morgan, Avner Snaked, and John Sutton, "Do people exploit their bargaining power? An experimental study," *Games and Economic Behavior*, 1991, *3* (3), 295–322.

Bohnet, Iris and Richard Zeckhauser, "Social comparisons in ultimatum bargaining," *The Scandinavian Journal of Economics*, 2004, *106* (3), 495–510.

Bolton, Gary E. and Axel Ockenfels, "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 2000, *90* (1), 166–93.

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, Gary Charness, and Matthew Ellman, "Let's talk: How communication affects contract design," Technical Report, CESifo Working Paper 2012.

Callen, Michael, Mohammad Isaqzadeh, James D Long, and Charles Sprenger, "Violence and risk preference: Experimental evidence from Afghanistan," *The American Economic Review*, 2014, *104* (1), 123–148.

Camerer, Colin and Teck Hua Ho, "Experience-weighted Attraction Learning in Normal Form Games," *Econometrica*, 1999, *67* (4), 827–874.

Cameron, Lisa Ann and Manisha Shah, "Risk-Taking Behavior in the Wake of Natural Disasters," *Journal of Human Resources*, forthcoming.

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.

Cohn, Alain, Ernst Fehr, Benedikt Herrmann, and Frdric Schneider, "Social Comparison and Effort Provision: Evidence from a Field Experiment," *Journal of the European Economic Association*, 2014, *12* (4), 877–898.

Dufwenberg, Martin and Georg Kirchsteiger, "A Theory of Sequential Reciprocity," *Games and Economic Behavior*, 2004, *47* (2), 268–298.

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 wagesthe 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*, forthcoming.

— , **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.

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.

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 John Moore**, "Contracts as Reference Points," *The Quarterly Journal of Economics*, 2008, *123* (1), 1119–11158.

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.

Kőszegi, Botond 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.

Peysakhovich, Alexander and David G. Rand, "Habits of Virtue: Creating Norms of Cooperation and Defection in the Laboratory," *working paper*, 2013.

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, 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.

Simonsohn, Uri, Niklas Karlsson, George Loewenstein, and Dan Ariely, "The tree of experience in the forest of information: Overweighing experienced relative to observed information," *Games and Economic Behavior*, 2008, *62* (1), 263–286.

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.

Online Appendix (Not for Publication)

A Results for Market Games

In this appendix we consider predictions for the PC and RC market games. Proposers' preferences (β, r_P) are drawn from some distribution F_P and responders' preferences (β, 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 the infimum of Proposer 1's offers, and let F be the distribution of those offers. Suppose that $\underline{a} < 1$. 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 chances of winning by at least $F(\underline{a} + \epsilon/2) - F(\underline{a})$. Of course, this can also lead to a loss in utility if an offer of size $a < \underline{a} + \epsilon$ would have been accepted. However, this loss is less than $(1 + \beta)\epsilon[F(\underline{a} + \epsilon/2) - F(\underline{a})]$, which is second order compared to a first-order gain from an additional increase in winning. Thus it is not possible to have $\underline{a} < 1$, which proves the proposition.

□

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 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 Alternative Models

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

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 MAOs than RC responders - the opposite of our experimental results. This 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 $a^{e_p} = 50$. 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

$$\int_x [a - \phi a(50 - a\tilde{\theta}]\mathbf{1}_{x \geq M} d\tilde{F}(x) \tag{8}$$

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

$$U_{2A} - U_{2R} = a + \phi\tilde{\theta}(300 - 3a - a)(100 - a) \tag{9}$$

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 Duwfenberg 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 (9) while it is increasing in equation (8)).

C Convergence results

We consider play in periods $t = -T, \ldots, 0, 1, 2, \ldots \infty$. In periods $t = -T, \ldots, 0$, players participate in some n player game (possibly one of the market games), while in periods $t = 1, 2, \ldots$ players participate in a non-competitive ultimatum game. As before, 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 $w_0, w_1, \ldots ..$ 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 market observation 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:

4. *Assume that $\gamma < 1$. Then there is a unique steady state $\langle(a^*, M^*), (r_P^*, r_R^*)\rangle$. In the steady state, $a^* > 0$, $a^* < k(W - a^*)$, 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:*

$$\lim_{t\to\infty} r_P^t = r_P^* \quad \textit{and} \quad \lim_{t\to\infty} r_R^t = r_R^*$$
$$\lim_{t\to\infty} a^t = a^* \quad \textit{and} \quad \lim_{t\to\infty} M^t = M^*$$

Proposition 4 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 4 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 final 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 MAOs, while RC responders should monotonically increase their MAOs:

5. *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 5 simply says that even though responders' MAOs should not reach steady state levels in a finite number of periods, the effect of past market experience should still diminish over time.

Proof. [Proof of Proposition 4]

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

$$M^* - \beta[r_R^*(k(Y - M^*) + M^*) - M^*] = 0, \tag{10}$$

which can be rearranged to show that

$$\frac{M^*}{k(Y - M^*) + M^*} = \frac{\beta r_R^*}{1 + \beta}. \tag{11}$$

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 (11), and using the definition of r_R^*, we now have that

$$r_R^* = (1 - \gamma)(1/2) + \gamma \frac{\beta}{1 + \beta} r_R^*. \tag{12}$$

Equation (12) is a linear equation in r_R^* with a unique solution given by

$$r_R^* = \frac{(1 - \gamma) + \beta(1 - \gamma)}{2 + 2\beta(1 - \gamma)}. \tag{13}$$

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 (13) 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

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

as $t \to \infty$. 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 **??**, 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 $|r_P^t + r_R^t - 1| \to 0$. 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

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

and

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

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

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 $\xi : \mathbb{R} \to \mathbb{R}$ as follows:

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

Define the map $\psi : \mathbb{R} \to \mathbb{R}$ 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 (11) implies that

$$\frac{M^*}{k(Y - M^*) + M^*} = \psi(\nu_R^*). \tag{14}$$

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 $\nu_R^t \to \nu_R^*$ will thus imply that $M^t \to M^*$ and $r_R^t \to r_R^*$. Moreover, since $|r_R^t + r_P^t - 1| \to 0$ 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 $t \geq t^\dagger$, $\nu_R^t \to \nu_R^*$ will thus also imply that $a^t \to a^*$.

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,

$$\nu_R^{t+1} = \frac{w_0}{\sum_{\tau=-T}^{t} w_{t-\tau}} \mu_R^t + \left(1 - \frac{w_0}{\sum_{\tau=-T}^{t} w_{t-\tau}}\right) \nu_R^t \tag{15}$$

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 $M^t \to M^{**}$. Thus

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

It is then easy to show that

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

On the other hand,

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

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. [Proof of Proposition 5]Since $r_P^t + r_R^t \leq 1$ for all $t \geq 1$, the reasoning of Step 3 in the proof of Proposition 4 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 4, 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.

□

D Comparing Phase 1 and Phase 2 proposer offers

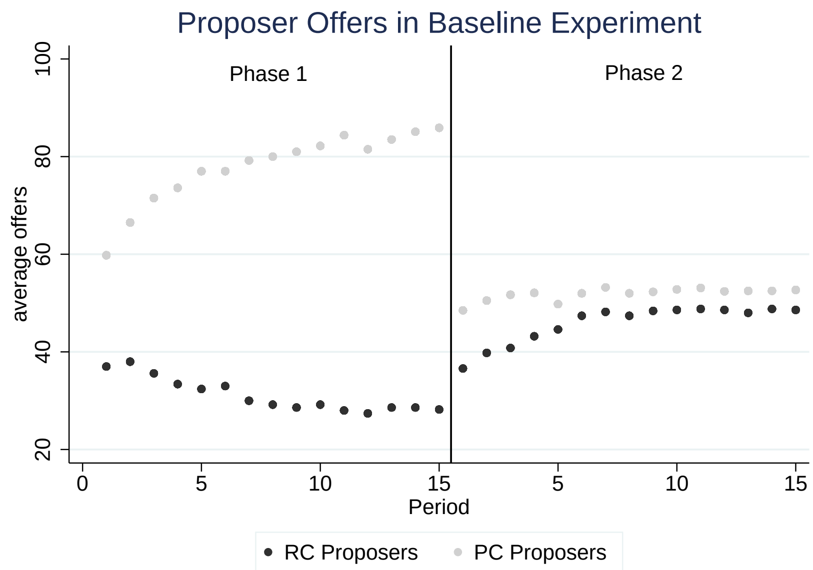

Figure 7: Proposer offers in Phase 1 and Phase 2 of the Baseline Experiment.

E Regressions with discounted experiences

In this section, we analyze whether our results on the effect of market observation and personal payoff experience are sensitive to their respective definitions. First, we have measured market observation 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, $G(\boldsymbol{\mu}^\tau, \boldsymbol{\nu}^\tau) = g\left(\sum \delta^t \mu^t / \sum \delta^t, \sum \delta^t \nu^t / \sum \delta^t\right)$, which give more weight to more recent periods in phase 1 of the experiment. Second, we have made assumptions about market observation in the Full Information treatment 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 10shows results of the same IV GMM regression that is presented in table 7 of the paper, but uses alternative measures of personal payoff experience and market observation.

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. Market observation and personal payoff experience remain significant determinants of responders minimal acceptance thresholds in phase 2 of the experiment. Also, the magnitude of the estimated coefficients remains quite stable. Consequently, our results on the

Table 10: Discounted Experiences

	(1)	(2)	(3)	(4)
market observation ($\delta = 0.9$)	23.843***	14.733***		
	(6.333)	(4.621)		
personal payoff experience ($\delta = 0.9$)	11.763**	10.120**		
	(5.143)	(5.018)		
market observation ($\delta = 0.95$)			24.852***	15.306***
			(6.670)	(4.843)
personal payoff experience ($\delta = 0.95$)			12.147**	10.461**
			(5.324)	(5.177)
Constant	32.031***	39.444***	31.891***	39.316***
	(4.743)	(3.150)	(4.787)	(3.169)
Adj. R^2	0.023	0.045	0.023	0.046
Observations	222	3330	222	3330

Market observation 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 7. Market observation and payoffexperience 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 competition group level (30 clusters). Significance levels: * = 10%, ** = 5% and *** = 1%.

impact of market observation 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 treatment, 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 market observation of subjects in the Full Information treatment.

Table 11 replicates table 7 from the main paper using these alternative codings of personal payoff experience and market observation. Personal payoff experience2 is the recoded

Table 11: Alternative codings of market observation and personal payoff experience

	(1)	(2)	(3)	(4)	(5)	(6)
market observation2	26.00***	15.96***			26.72***	16.43***
	(7.06)	(5.11)			(6.93)	(4.97)
personal payoff experience	12.59**	10.84**				
	(5.52)	(5.36)				
market observation			26.72***	16.43***		
			(6.93)	(4.97)		
personal payoff experience2			12.40**	10.92**	12.41**	10.92**
			(5.48)	(5.34)	(5.48)	(5.33)
Constant	31.71***	39.16***	31.26***	38.83***	31.26***	38.83***
	(4.84)	(3.19)	(4.74)	(3.15)	(4.74)	(3.15)
Adj. R^2	0.025	0.047	0.030	0.050	0.030	0.051
Observations	222	3330	222	3330	222	3330

market observation2 differs from market observation in how observed offers in the Full Information treatment are weighted. market observation 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. market observation2 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 X. market observation 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 competition group level (30 clusters). Significance levels: * = 10%, ** = 5% and *** = 1%.

personal payoff experience variable and market observation2 is the recoded market observation variable. Columns (1)-(6) replicate the original estimations in table 7 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 market observation and personal payoff experience. The magnitude of the estimated coefficients is remarkably stable, and market observation and personal payoff experience remain significant determinants of responders' minimal acceptance thresholds independent of the precise definition of these terms.

F Ruling out Anchoring

As we note in section 4.2 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.[35] 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 affect 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 treatment. In this treatment, 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 responder's acceptance behavior were influenced by the provision of arbitrary anchors, we should observe that responders in the RC market in the Full Information treatment show higher acceptance rates than Responders in the RC market in the Baseline or in the Role Switch treatment.[36] This is the case because in the Full Information treatment, they are subjected to higher anchors than in the Baseline treatment or in the Role Switch treatment, in which responders only get to observe the offers made by their matched proposer.

Table 12 contains information that tests whether the information provided in the Full information treatment indeed provides an alternative anchor and consequently changes proposers' offers. Moreover, table 13 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 behavior

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

[36]We 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 12: Proposer offers

	(1)	(2)
full info * propcomp	1.316	1.758
	(0.982)	(1.395)
full info * respcomp	–3.732	2.106
	(3.606)	(3.766)
propcomp	41.337***	46.733***
	(3.071)	(3.396)
Constant	36.985***	31.147***
	(2.991)	(3.173)
Adj. R^2	0.689	0.719
Observations	3330	2430

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

of responders in the Full Information treatment is statistically different from the respective behavior in the Baseline or the Role Switch treatment. 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 treatment. Consequently, anchoring does not seem to be a driving force of behavior in our experimental setting.

Table 13: Responder Acceptance Decisions

	(1)	(2)
maxoffer	0.006***	0.006***
	(0.001)	(0.001)
full info * propcomp	–0.051	–0.014
	(0.048)	(0.041)
full info * respcomp	–0.004	–0.038
	(0.029)	(0.044)
propcomp	0.019	–0.044
	(0.043)	(0.084)
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. propcomp indicates observations from the Proposer Competition market. full info*propcomp is an interaction between a Full Information treatment dummy and a proposer competition market dummy. Equivalently, full info*respcomp is an interaction between a Full Information treatment dummy and a responder competition market dummy. Best offer 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 treatment and the Baseline treatment. Standard errors are clustered at the phase 1 competition group level (30 clusters in columns (1) and (3), 22 clusters in columns (2) and (4). Significance levels: * = 10%, ** = 5% and *** = 1%.

G The Full Information Experiment

Table 14: Minimum Acceptance Thresholds in the Full Information Experiment

	OLS		IV	
	(1)	(2)	(3)	(4)
PCResponder	11.552*	5.644		
	(6.106)	(5.644)		
payoff experience			19.694 **	9.622
			(9.974)	(9.272)
Constant	48.448***	47.575***	47.171 ***	46.951 ***
	(4.674)	(3.897)	(4.807)	(3.983)
Adj. R^2	0.013	0.034		
Observations	87	1305	87	1305

Columns (1) and (2) show results of an OLS regression. Proposer phase 1 market experience dummy included. Column (1) uses data from period 1 of phase 2 only. Column (2) uses all data. Columns (3) and (4) show IV GMM regressions, payoff experience is instrumented using 6 dummies, one for each market treatment and 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 competition group level (18 clusters). Significance levels: * = 10%, ** = 5% and *** = 1%.

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

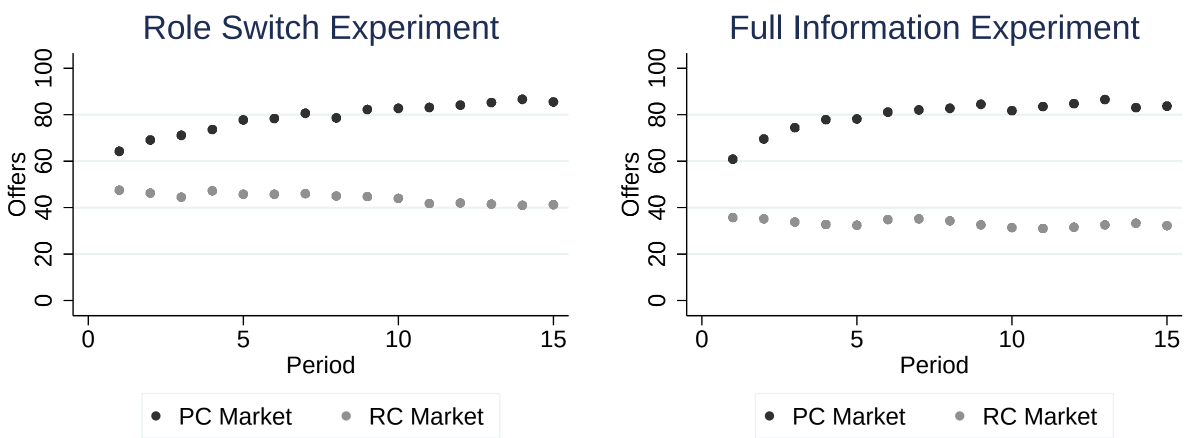

Figure 8: Left Panel: Average offers over time under responder competition and under proposer competition in phase 1 of the Role Switch experiment. Right Panel: Average offers over time under responder competition and under proposer competition in phase 1 of the Full Information experiment

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 ($p < 0.001$).