UNIVERSITY OF CAMBRIDGE
Faculty of Economics

Institute for New Economic Thinking

Cambridge-INET Institute

Cambridge-INET Working Paper Series No: 2017/20

Cambridge Working Papers in Economics: 1743

COMMITMENT AND (IN)EFFICIENCY: A BARGAINING EXPERIMENT

Marina Agranov
(California Institute of Technology)

Matt Elliott
(University of Cambridge)

We conduct an experimental investigation of decentralized bargaining over the terms of trade in matching markets. We study if/when efficient matches are reached, and the terms of trade agreed upon. We find that mismatch is extensive, and persists as we change the nature of bargaining by moving from a structured experimental protocol to permitting free-form negotiations. We identify two sources of inefficiencies. Inefficiencies are driven by (a) players' rational responses to their bargaining positions changing as others reach agreement, and (b) the existence of players who are unwilling to accept low, inequitable payoffs.

Commitment and (In)Efficiency: a Bargaining Experiment*

Marina Agranov[†] Matt Elliott[‡]

September 3, 2017

Abstract

We conduct an experimental investigation of decentralized bargaining over the terms of trade in matching markets. We study if/when efficient matches are reached, and the terms of trade agreed upon. We find that mismatch is extensive, and persists as we change the nature of bargaining by moving from a structured experimental protocol to permitting free-form negotiations. We identify two sources of inefficiencies. Inefficiencies are driven by (a) players' rational responses to their bargaining positions changing as others reach agreement, and (b) the existence of players who are unwilling to accept low, inequitable payoffs.

1 Introduction

A fundamental question in economics is whether the "right" people end up in the "right" jobs. Labor markets are important and their allocative efficiency is crucial for the productivity of the economy. The typical way that frictions have been incorporated into these markets is through costly search and imperfect information.[1]

*We thank Federico Echenique, Edo Gallo, Ben Gillen, Ben Golub, Francesco Nava, Emanuel Vespa, Leeat Yariv, and participants in seminars at Berkeley HAAS, Berlin, Bonn, Cologne, Columbia, EUI, Georgetown, NYU, Pittsburgh, Queen Mary, Tel Aviv, UCL, and UCSD for helpful comments and suggestions.

†Division of the Humanities and Social Sciences, California Institute of Technology, 1200 E. California Blvd, Pasadena, CA 91125 USA. Email: magranov@hss.caltech.edu.

‡Cambridge University, Faculty of Economics. Email: mle30@cam.ac.uk.

[1] See Rogerson et al. (2005) and Rogerson and Shimer (2011) for surveys of the search literature. Informational frictions are studied in Calvo-Armengol and Jackson (2004), among others.

However, in many high-skill labor markets these frictions are limited. Workers typically know which firms are looking to hire, and similarly firms know which workers would be appropriate for a given vacancy. Does this mean that the right workers will end up in the right firms? More specifically, absent these frictions can decentralized bargaining be expected to result in an efficient allocation of workers to firms?

On the one hand, the Coase Theorem suggests that as long as transaction costs are sufficiently small, two parties should be able to bargain to an efficient outcome. This motivates the idea that no worker-firm pair should leave gains from trade on the table, and this is well known to be a sufficient condition for matching markets to clear efficiently (Shapley and Shubik, 1972).

On the other hand, agreements in decentralized matching markets are typically reached sequentially. Thus the composition of the market, those workers and firms who are actively searching, changes and as this market context evolves so can the bargaining positions of the remaining workers and firms. Suppose it is efficient for a worker, Ann, to match to a firm, B, but that Ann is currently in a strong negotiating position; There is another firm, C, with a vacancy Ann could instead fill. Although Ann would be less productive with firm C, Ann would still like to use this alternative possible match to bid up her wage with firm B. However, this alternative vacancy at C might be filled by someone else, in which case Ann would lose her strong bargaining position. Indeed, if there is no chance Ann will match inefficiently with firm C, then firm B might as well wait for C's vacancy to be filled and for Ann's bargaining position to deteriorate. Can agents who find themselves in temporarily strong bargaining positions benefit from these positions without sometimes matching inefficiently?

No empirical work we are aware of investigates whether bargaining frictions, i.e., the strategic actions of market participants to improve their terms of trade, can lead to allocative inefficiency. Fundamental identification problems inhibit such an investigation. Even under very strong assumptions it is hard to identify whether matches are positively or negatively assortative from wage data (Eeckhout and Kircher, 2011).[2] More generally, to observe the extent of mismatch, an econometrician must estimate the counterfactual productivities of matching different people to different jobs. But unobservable worker characteristics that are valued differently by different firms can generate any counterfactual productivities and rationalize any given match as efficient. Even were it possible to observe detect match inefficiencies, it would be hard to separate out the role of bargaining frictions with other frictions. To overcome these problems, and provide some first empirical evidence on bargaining frictions, as opposed to search or informational frictions, we take an experimental approach.

We use an array of laboratory experiments to study how payoffs are affected by the

[2]Specifically, assuming that workers are vertically differentiated, firms are vertically differentiated, and surpluses are either supermodular or submodular. Under these assumptions the welfare loss from inefficient matching *can* be identified.

structure of the market and whether allocative efficiency, matching the "right" worker to the "right" job, is achieved by decentralized bargaining. Matching is one-to-one and we consider the simplest markets in which a player can lose a matching possibility as others reach agreement and exit. Our experiments feature two characteristics which are common in many labor markets: heterogeneous match surpluses and endogenous agreements regarding how the surplus generated by a match is split.[3] In the lab, we control the entire set of possible match surpluses, removing unobserved heterogeneity in match quality and observing counterfactual match productivities. We also track individuals' bargaining patterns in full.

Our study comprises three different experiments (Experiment I, II, and III), which differ in the details of bargaining protocol as we describe below. Each experiment consists of several treatments, corresponding to different market structures.

In Experiment I, we use a standard bargaining protocol that mirrors one often used in the theoretical literature to study three simple markets. We find that the market composition at the time an agreement is reached affects the average payoffs of our experimental subjects—market composition affects bargaining positions. Moreover, workers and firms frequently match inefficiently, and we show that the extent of mismatch is higher when players' bargaining positions change more as others exit the market.

In practice, however, interactions in markets are much less structured that those imposed by this first experimental protocol, or in fact any protocol corresponding to a dynamic bargaining model from the theoretical literature. Moreover, it might be expected that in richer environments, one way or another, the Coase Theorem will prevail and bargaining will lead to an efficient outcome. For example, efficiency enhancing norms might arise in these less structured environments, or alternatively, protocol-free bargaining could be more conducive to more complicated but efficient play. We investigate this possibility in Experiment II by permitting unstructured interactions between the market participants. The market composition at the time agreements are reached continues to affect the terms of trade, and despite the above intuition we find that even when bargaining is protocol-free inefficiencies persist.

Finally, in Experiment III, we investigate further the mechanism that drives the inefficiencies we observe in the first two experiments. Data from Experiments I and II suggests a strong connection between inefficient matching and non-stationarities in bargaining positions. Building on this, in Experiment III we directly limit non-stationarities by allowing players to renege on agreements they have already reached

[3]Many markets are characterized by heterogeneous surpluses, and so it matters which worker is employed by which firm. We expect, however, that getting the "right" worker to be employed by the "right" firm is likely to matter more in high-skill labor markets. These markets are also characterized by wages being negotiated (Hall and Krueger, 2012) and are economically important: the top 10% of earners accounted for 45% of overall income and for 68% of federal income tax receipts in the US in 2011 (`http://www.heritage.org`).

for a small cost, while keeping everything else the same as Experiment I. We find that, consistent with the aforementioned intuition, the option of reneging substantially increases efficiency, although non-trivial rates of mismatch remain.

A more in-depth analysis of the data generated by Experiment I suggests that there are two major sources of inefficiency. This first is players' rational responses to their bargaining positions changing as others reach agreements, while the second is that some players are unwilling to accept low and inequitable payoffs when it would be rational to do so. One way to interpret our results from Experiments II is that the unstructured nature of interactions helps remove inefficiency due to players rejecting low offers it would be rational to accept, but inefficiency due to players' rational responses to their non-stationary bargaining positions remains. While the results from Experiment III can be interpreted as suggesting that inefficiencies due to changing bargaining positions are removed by permitting reneging, but not the inefficiencies due to players rejecting low offers they should accept.

We contend that in many real labor markets the bargaining positions of players change as others reach agreements and exit the market. Our experimental investigation replicates and studies this feature. Alternative matches affect the average terms of trade agreed upon. The composition of the market, which workers and which firms are still searching for a match, thus matters and players' bargaining positions are non-stationary. Evidence across our three experiments collectively suggests that this non-stationarity is intimately tied to high rates of inefficient matching. Our experiments provide some first evidence for the role of bargaining frictions, as opposed to search or informational frictions, in decentralized matching markets.

1.1 Related Literature

We focus in this section on the related experimental literature. We discuss the theoretical literature in the context of our different experimental protocols.

There is a large experimental literature on bargaining.[4] The most relevant to our paper is the study by Binmore et al. (1989), which investigates the effect of exogenous outside options on the bargaining position of players in a two-person bargaining setup that has features of both the alternating-offer and ultimatum-game protocols. The authors find that responders receive a payoff equal to their binding outside option, providing support for the "outside option principle."

Our paper is also related to the experimental literature on decentralized two-

[4]See Roth (1987) for an overview of experimental work on coalition bargaining, which was mostly concerned with testing cooperative game theory concepts, Roth (1995) for a survey of early experiments exploring non-cooperative theories of bargaining, and Palfrey (2016) for a recent survey of multilateral bargaining games.

sided matching markets, which is relatively thin and for the most part focuses on matching markets with non-transferable utility. The prominent studies in this space include Echenique and Yariv (2013) and Pais and Veszteg (2011). Echenique and Yariv (2013) consider fully decentralized two-sided matching markets with complete information and find that most markets reach stable outcomes. When more than one stable outcome exists, the outcomes gravitate towards the median stable match. Pais and Veszteg (2011) study both complete and incomplete information matching markets and vary search costs and the degree of commitment to formed matches; this last feature is reminiscent of the variation in the bargaining protocol we consider in this paper. The authors find that in complete information markets, which are the closest to our setup, the degree of commitment affects both the frequency of efficient final matchings and the level of market activity as captured by the number of match offers made by subjects. Contrary to our main finding, the authors document that the treatments with commitment correspond to the highest proportion of efficient final outcomes.[5]

The main difference between our paper and those discussed above is that we allow bargaining over the terms of trade, studying decentralized matching markets with *transferable utility*. This brings to light an additional dimension of the bargaining process which is missing, by construction, from games with non-transferable utility: Bargainers need to agree not only on who is matched to whom but also how to split the available surplus between the pair of potential match partners. The only other experimental study of decentralized matching markets with transferable utility that we are aware of is the study by Nalbantian and Schotter (1995). In this paper, the authors analyze several procedures for matching with players who are privately informed about their payoffs.[6] The authors find that while efficiency levels were relatively high in all treatments, different mechanisms suffer from different types of problems: Some produce a considerable number of no-matches while others produce a substantial number of suboptimal matches.

There is a small experimental literature studying bargaining on networks, which is surveyed in Choi et al. (2016). The study most closely related to ours is Charness et al. (2007), which examines experimentally the effects of network structure on market outcomes following the model of Corominas-Bosch (2004). The bargaining is structured as a sequential alternating public-offer bargaining game over the shrinking

[5]For an experimental study of one-sided matching markets with non-transferable utility see Molis (2010). For studies with a more rigid bargaining structure, such as the one in which one side of the market makes offers to the other side but not vice-versa, see Haruvy and Ünver (2007) and Niederle and Roth (2009). Finally, see Kagel (2000) and Featherstone and Mayefsky (2010), who study unravelling and the transition between a decentralized market and a centralized clearinghouse.

[6]These matching procedures range from the free-agency system similar to the problem of matching baseball players to teams, to the simultaneous bid mechanism, in which participants on each side of the market simultaneously submit the maximum amount they are willing to pay to be matched to each participant on the other side of the market.

value of homogeneous and indivisible goods. Offers made by players on one side of the market alternate with offers made by players on the other side of the market, and all players on a given side of the market make offers simultaneously. An offer is a price which is announced to all players on the other side of the market, who then choose which offers to accept. Experimental results qualitatively support the theoretical predictions and display a high degree of efficiency: Total payoffs of players constitute over 95% of the maximum attainable surplus, and three-quarters of all agreements are reached in the first bargaining round.[7]

Finally, there is an experimental literature in sociology that studies how network structures confer power. Two foundational papers are Cook and Emerson (1978) and Cook et al. (1983), and there is a nice albeit brief discussion in Jackson (2010). A typical experimental design in this literature has several features different from us, and more importantly, the focus is on identifying strong and weak network positions rather than evaluating the efficiency of markets. As far as we are aware, this literature does not investigate the interaction between changing market composition and efficiency, and the typical protocol considered does not lend itself to such an investigation by preventing players from exiting before negotiations among all possible matches have taken place.[8]

2 Environment

We set out to test whether the endogenous evolution of thin, heterogeneous matching markets can result in an inefficient allocation of workers to firms in the case of labor markets, buyers to sellers in product markets, or men to women in the marriage market. As we suspect that inefficiencies will be more likely in more complicated settings, we consider the simplest possible markets capable of exhibiting the effects we are interested in. For bargaining positions to change as others exit we need the market to be able to support at least two matches, requiring at least four players and at least three different matches among these four players to be possible.

Figure 1 presents three different market structures (Game 15, Game 25, and Game 30) which will serve as a basis of our investigation. These are four-person markets, with each player identified by the letter A, B, C, or D. A link between two players indicates the joint surplus that this pair of players generate by matching with each other, with the surplus indicated by a number next to the link. These are one-to-one matching markets, that is, each player can be matched with at most one other player in the market. The payoffs of unmatched players are normalized to 0. In all three

[7]See also Gale and Kariv (2009) and Choi et al. (2014) for a study of trading in networks with intermediaries, implemented through a simultaneous bid-ask protocol and posted prices respectively.

[8]Perhaps closest papers to ours are Bienenstock and Bonacich (1993) and Skvoretz and Willer (1993). Both consider a variety of theories including some from cooperative game theory.

Figure 1: The three markets considered in this study.

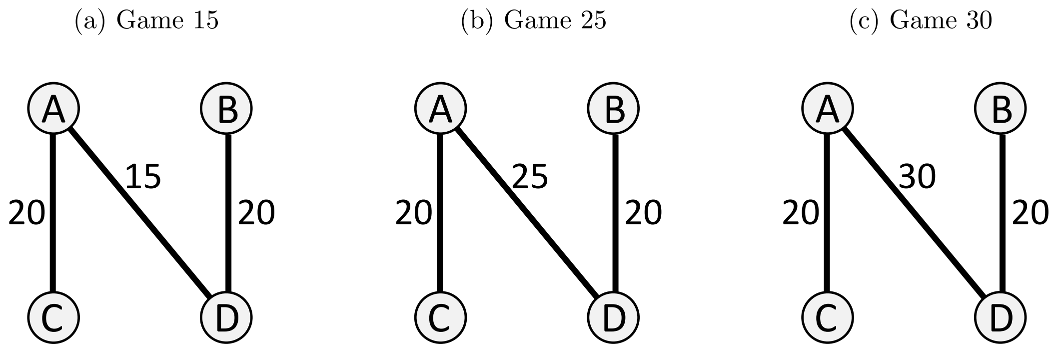

Notes: We refer to players A and D as the strong players and players B and C as the weak players.

markets, the vertical links (the link between A and C and the link between B and D) generate a surplus of 20 units. The markets differ in one feature only: the value of the diagonal link between A and D. In Game 15 this link is worth 15 units, in Game 25 it is worth 25 units, and in Game 30 it is worth 30 units. This diagonal link determines the bargaining position of A and D vis-a-vis C and B. We will refer to A and D as the strong players, and to B and C as the weak players. In all three markets it is efficient for A and C to match and for B and D to match. Across these three markets we study how the average payoffs of players differ and the frequency with which the efficient match is reached. A key question we investigate is whether A and D can use the alternative possibility of matching with each other to secure a larger share of the surplus when matching to their efficient partners, without having to actually match inefficiently. Is the threat of matching inefficiently enough to increase A and D's expected payoffs to the point where they no longer want to match inefficiently, or does the inefficient match have to be exercised to be credible?

The general version of the problem we study is how a finite set of workers W are allocated among a finite set of firms F, each of which have a single vacancy. Normalizing the value of remaining unmatched to 0 for all players, we let s_{ij} be the surplus that worker i would generate with firm j. A bargaining outcome is a tuple $(\mathbf{u}, \mathbf{v}, \mu)$ where $\mathbf{u}$ is a vector of worker payoffs, $\mathbf{v}$ are the firms' payoffs and, letting $N = W \cup F$, the match $\mu : N \to N$ describes which worker is matched to which firm.[9] There is generically a unique match that maximizes the total surplus generated in the market.[10] We denote this efficient match by $\mu^* : N \to N$, so that

[9]This function must satisfy the usual restriction that i is matched to j ($\mu(i) = j$) if and only if j is matched to i ($\mu(j) = i$). Following convention, we let $\mu(i) = i$ represent that i is unmatched. A worker i must then be matched to either a firm or herself, while a firm j must be matched to a worker or himself.

[10]If independent error terms were drawn from a continuous, atomless distribution and added to each surplus, the efficient match would be unique with probability 1.

$$\sum_{\text{workers } i} s_{i\mu^*(i)} = \max_{\mu} \sum_{\text{workers } i} s_{i\mu(i)}.$$

A major focus of our analysis will be the probability with which our experimental subjects match efficiently. In Game 15, Game 25 and Game 30, this is the probability that A matches to C and B matches to D.

3 Experiment I

3.1 Experimental Design and Procedures

Experiment I consists of three treatments (BASELINE 15, BASELINE 25, and BASELINE 30) corresponding to the three markets described in Figure 1. All our experimental sessions were conducted at two locations: the Experimental Social Science Laboratory (ESSL) at University of California, Irvine and the Experimental and Behavioral Economics Laboratory (EBEL) at University of California, Santa Barbara. At both locations, subjects were recruited from a database of undergraduate students enrolled in these universities.[11,12] Ten sessions were conducted, with a total of 172 subjects. No subject participated in more than one session. The experiments lasted about one hour and a half. Average earnings, including a \$15 show up fee, were \$23.5 with a standard deviation of \$5.3.

In each experimental session subjects played ten repetitions of the same game with one or more rounds in each repetition and random re-matching between games (i.e. between repetitions). In other words, before the beginning of each game subjects were randomly divided into groups of four and assigned one of the four letters (A, B, C or D), which determined their network position. This procedure is standard practice in the experimental literature and is often used in relatively complicated games in which it is natural to expect learning.

Within each game, we implement the following bargaining protocol. At the beginning of a game all players are unmatched. At the beginning of each round *all* unmatched players then choose a) whom, if anyone, to make an offer to and b) how to split the available surplus. One player is then selected at random to be the proposer, and her offer is implemented. This timing differs in a strategically irrelevant way from the game described and allows us to collect more data on proposals. If

[11] The software for the experiment was developed from the open source Multistage package, available for download at http://software.ssel.caltech.edu/.

[12] In the Supplementary Appendix, Section 1.3, we report the location at which each session was conducted and compare the behavior of subjects across the two labs. Our data suggest that there is no significant difference between the subject pool at UCI and the one at UCSB.

the offer of the selected player is rejected, then both players remain unmatched and the group proceeds to the next round of the game. If the offer is accepted, then the matched players exit the market permanently. All players in the group observe the move of the selected player and the move of the responder. There are two ways in which the game can come to an end. The first one is the situation in which the surplus generated by any pair of unmatched players who have made proposals in the last round is 0. The second one is discounting implemented as a random termination of the game: There is a 1% chance that each round is the last one in a game and a 99% chance that the game is not over. When the game ends unmatched players receive a payoff of 0 while matched players earn payoffs according to their agreements. At the end of the experiment, the computer randomly selects one of the ten games played, with all ten game being equally likely to be selected. Subjects' earnings in the experiment consist of a show-up fee plus their earnings in the randomly selected game.

In the Supplementary Appendix, Sections 6.1, we present the instructions that were distributed to the subjects and read out loud by the experimenter before the beginning of the experiment. Before starting the experiment, subjects were asked to complete a quiz, which tested their understanding of the game rules. Subjects could not move on to the experiment until they correctly answered all the questions on the quiz.[13] Two features of our interface are worth mentioning. First, at all times the subjects saw the network structure and the available surpluses on the left-hand side of the screen. Second, on the right-hand side of the screen, subjects could observe how the matches evolved over the course of the previous rounds for the current game, by clicking arrow buttons below the diagram that depicted the network structure. These features were implemented to ensure that the subjects had complete information about what had transpired in the previous rounds of a game, in order to eliminate reliance on the subjects' memory of the history of play.

3.2 Experiment I: Results

We will consider efficiency levels, payoffs and strategies across the three markets.

3.2.1 Approach to Data Analysis

Our main interest is in subjects' behavior after they have had the opportunity to experience the game. Allowing for the presence of an initial learning phase, our statistical tests use data from the last five repetitions of a game played in each experimental session. We think play in these games will better reflect the market situations we

[13]The list of questions and the screenshots of the game are presented in the Supplementary Appendix, Sections 6.2 and 6.3.

seek to capture—in these markets the stakes are higher and many participants will have some experience. We refer to these games as *experienced games*. We refer the reader to the Supplementary Appendix, Section 2 for the detailed analysis of initial repetitions of the game and learning.

When we compare outcomes between two groups, we focus on the *average* outcomes. Sometimes the two groups will correspond to two different markets. At other times we will fix the market under consideration and compare the payoffs of players in strong bargaining positions with those in weak bargaining positions. To compare the outcomes of two groups we run random-effects GLS regressions in which we regress the variable of interest (i.e., the payoffs of the players or an indicator of whether the final match is efficient) on a constant and on an indicator for one of the two groups considered. To account for interdependencies between observations that come from the same session, we cluster standard errors by session. We say there is a significant difference between the outcomes of the two groups under consideration if the estimated coefficient on the group indicator dummy is significantly different from 0 at the 5% level, and we report p-values associated with that estimated coefficient.

When we compare final outcomes between games, we focus on the groups that finished the game naturally rather than those that were interrupted by the random termination.[14] When we investigate market dynamics and the strategies used by our experimental subjects, we use all the collected data, including all the submitted proposals rather than just proposals randomly selected for implementation.

3.2.2 Market Outcomes

Table 1 summarizes the outcomes in each market.[15] We observe a significant decrease in efficiency levels as the value of the inefficient match (the diagonal link) increases from Game 15 to Game 25 to Game 30. Indeed, while all the final matches in BASELINE 15 treatment, in the experienced games, are efficient, the probability the efficient match is reached drops to 51% in BASELINE 25 and even further to 30% in BASELINE 30 treatments. Regression analysis confirms that this monotonic decrease in efficiency is significant ($p < 0.01$ for BASELINE 15 vs. BASELINE 25, and $p = 0.01$ for BASELINE 25 vs. BASELINE 30).

Further, we find that the network position of the first mover affects the likelihood of reaching the efficient match in Games 25 and 30. Recall that we refer to players with one link (players B and C) as the weak players, and those with two links (players A and D) as the strong players. Our data shows that the efficient matching is more likely to be reached in BASELINE 25 and BASELINE 30 treatments when the weak

[14] Random termination was very rare: about 6% of games in all treatments of Experiment I ended because of random termination.

[15] The evolution of final match efficiency is presented Section 1.1 in the Supplementary Appendix.

players propose first. To detect this pattern, for each of the two treatments separately (BASELINE 25 and BASELINE 30), we regress an indicator variable for the final match being efficient on an indicator variable for the first mover being a weak player. In both BASELINE 25 and BASELINE 30 treatments we find that efficiency is significantly lower in games in which the first mover is a strong player rather than a weak player ($p < 0.05$). The same is true if we condition the efficiency of the final match on the network position of the player who makes the first accepted offer. Details of this analysis are presented in Section 1.2 in the Supplementary Appendix. We summarize our analysis of efficiency in Result 1.

Result 1.

(i) Matching is efficient in BASELINE 15 treatment.

(ii) The rate of efficient matching declines monotonically from BASELINE 15 to BASELINE 25 and then to BASELINE 30.

(iii) In BASELINE 25 and BASELINE 30 treatments efficient outcomes are more likely to be reached if a weak player is selected to propose first.

Table 1: Observed outcomes in Experiment I, experienced games

	BASELINE 15 treatment			BASELINE 25 treatment			BASELINE 30 treatment		
	eff.	B (C)	A (D)	eff.	B (C)	A (D)	eff.	B (C)	A (D)
Observed	100%	10 (0.05)	10 (0.05)	51%	4.5 (0.36)	11.8 (0.13)	30%	2.4 (0.32)	14.2 (0.16)

Notes: The rate at which the efficient match is reached and the average payoffs of players, by their network position, are reported with the corresponding robust standard errors in the parenthesis where observations are clustered at the session level.

We now turn our attention to players' payoffs, distinguishing between payoffs of weak players (players B and C) and payoffs of strong players (players A and D). Table 1 reports the average payoffs of the players, along with the corresponding standard errors.

In BASELINE 15 treatment players' payoffs do not depend on their network position ($p = 0.319$), while in both BASELINE 25 and BASELINE 30 treatments strong players receive significantly higher payoffs than weak players ($p < 0.01$). The value of the diagonal links affects the bargaining position of strong players insofar as as they obtain significantly higher payoffs in BASELINE 30 than in BASELINE 25, and in BASELINE 25 than in BASELINE 15. At the same time weak players obtain significantly lower payoffs in BASELINE 30 than in BASELINE 25, and in BASELINE 25 than in BASELINE 15 ($p < 0.01$ in all pairwise comparisons).

Figure 2: Final payoffs of the strong players in efficient matches in Experiment I by market composition at the time of exit, experienced games

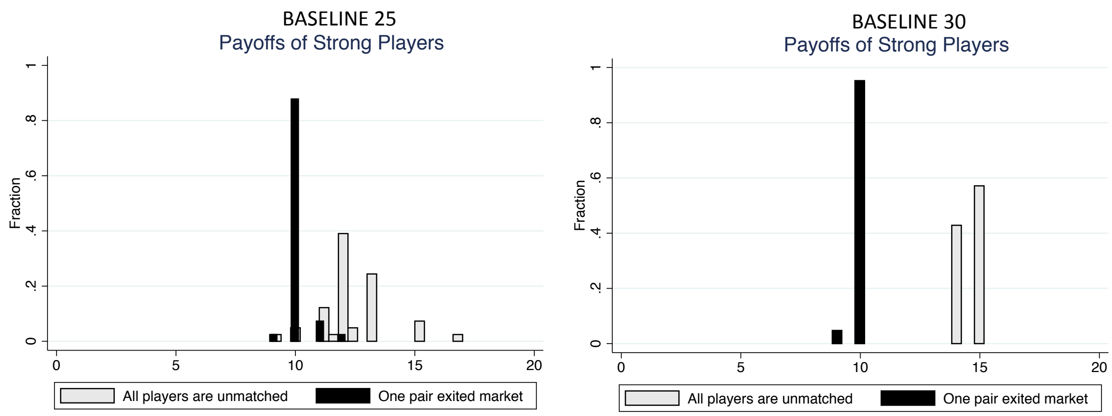

Notes: In both figures, we consider only groups that reached efficient match.

To further examine the affect of the diagonal link on bargaining positions, we consider the payoffs of the strong and weak players when they exit first versus second, conditional on the final match being efficient. In contrast to when players exit the market first, when they exit the market second, they are engaged in bilateral bargaining and they have identical bargaining positions. In BASELINE 25 treatment, the average payoff of strong players is 12.3 if they exited first, while it is only 10.1 if they exited second. For weak players it is 7.7 if they exited first and 9.9 if they exited second.[16] Similarly, in BASELINE 30 treatment, the average payoff of strong players is 14.5 if they exited first, while it is 10 if they exited second; the weak players' average payoff is 5.5 when exiting first compared to 10 when exiting second. Regression analysis confirms that these differences are statistically significant with $p < 0.01$ in both treatments.[17] Regression analysis also confirms that this difference is increasing in the value of the diagonal link. The difference in strong players' payoffs when matching first instead of second is higher in BASELINE 30 than BASELINE 25, and in BASELINE 25 than BASELINE 15.[18] We obtain similar results when

[16]This does not incentivize weak players to delay in equilibrium because doing so increases the probability they will be left unmatched.

[17]For comparison, when the inefficient match is reached weak players receive a payoff of 0 and the average payoff of a strong player is 12.5 in BASELINE 25 and 15 in BASELINE 30.

[18]Statistical analysis of this claim is based on several pairwise comparisons. First, we establish that payoffs of strong players who exit market second are not statistically different in all three games ($p > 0.10$ in all three pairwise comparisons). Second, we find that payoffs of strong players who exit market first monotonically increase as we move from BASELINE 15 to BASELINE 25 to BASELINE 30 ($p < 0.01$ for both BASELINE 15 versus BASELINE 25 and BASELINE 25 versus BASELINE 30).

using two observations per subject: average payoff of a subject when she was in a strong position and exited first and average payoff of a subject when she was in a strong position and exited second.[19] Illustrating this, Figure 2 presents histograms of the final payoffs of the strong players by their order of exit and conditional on an efficient outcome being reached, where grey bars depict payoffs of players that exited the market first and black bars represent payoffs of players that exited the market second.[20]

We summarize our analysis of payoffs in Result 2.

Result 2. *Moving from BASELINE 15, to BASELINE 25, to BASELINE 30*

(i) The payoffs of strong players monotonically increase.

(ii) The payoffs of weak players monotonically decrease.

(iii) Conditional on reaching the efficient final match, the difference in the payoffs of the strong players from exiting first rather than second is positive and monotonically increases.

An immediate implication of Result 2 (iii), is that the opposite pattern holds for weak players: Conditional on reaching the efficient final match, the difference in the payoffs of the weak players from exiting second rather than first is positive and monotonically increases from BASELINE 15 to BASELINE 25 and from BASELINE 25 to BASELINE 30.

Finally, we consider players' strategies. Players rarely choose the "Do Nothing" button (less than 2.5% in both BASELINE 25 and BASELINE 30 treatments), which can be seen as the manifestation of delay, i.e., the choice not to make a proposal. We observe no significant difference between the frequency of delays between the strong and the weak players conditional on the market being complete (such that no players have exited) with $p > 0.10$ in both BASELINE 25 and BASELINE 30 treatments.

[19] This refutes the concern that this result is driven by the selection of subjects, e.g., that some subjects are better at bargaining so tend to obtain higher payoffs, and these subjects also tend to exit the market first when in strong positions. Specifically, using two observations per subject, in BASELINE 25 treatment, the average payoff of strong players when exiting first is 12.0, while the average payoff of strong players when exiting second is 10.1. Similarly, in BASELINE 30 treatment, the average payoff of strong players when exiting first is 12.9, while the average payoff of strong players when exiting second is 10.1. Moreover, while different subjects have different numbers of times that they were assigned to the position of a strong player and exited first or second, in BASELINE 25 treatment, for 65% (84%) of subjects, the number of times they exited first versus second differs at most by one (two) instance(s). The same statistics for BASELINE 30 treatment are 77% and 98%, respectively.

[20] Only the payoffs of strong players are shown, but the expected payoffs of weak players conditional on reaching the efficient match are 20 less the payoff of the strong players conditional on reaching the efficient match.

When all players are present in the market strong players also have to decide whether to offer to each other or to their efficient partners. We observe a significant shift in the strategies used by the strong players as the value of the diagonal link increases. In particular, in BASELINE 15 treatment over 80% of subjects always propose only to their efficient partner in every game in which they are assigned to a strong bargaining position. In contrast, in BASELINE 30 treatment over 80% of subjects never propose to their efficient partner when in strong bargaining positions. To statistically examine this pattern, we focus on the proposals made by strong players and regress an indicator variable for whether the proposal was made to the subject's efficient partner, on an indicator variable for the game. Regression analysis confirms that subjects in strong network positions are more likely to make efficient proposals in BASELINE 15 than in BASELINE 25 ($p < 0.01$), and in BASELINE 25 than in BASELINE 30 ($p < 0.01$).

We summarize our analysis of strategies in Result 3.

Result 3.

(i) Players do not delay.

(ii) The frequency with which the strong players make efficient offers monotonically declines from BASELINE 15, to BASELINE 25, to BASELINE 30.

3.3 Theoretical Predictions

To guide the interpretation of our experimental results it is helpful to consider some alternative theories. These theories yield different predictions of players' expected payoffs and matches—and thus the level of efficiency in the market. For each theory, we briefly describe the main idea and implications for the three games depicted in Figure 1. We refer the reader to Sections A and B in the Appendix for additional details.

3.4 Cooperative Theory

According to the Coase Theorem, as long as transaction costs are sufficiently low, decentralized bargaining should result in an efficient outcome. This provides strong foundations for expecting decentralized bargaining in matching markets to result in outcomes that are robust to pairwise deviations, such that there is no buyer-seller pair who could both do better by reaching some agreement between themselves. Shapley and Shubik (1972) show that ruling our pairwise deviations in matching environments such as ours is necessary and sufficient for ruling our coalitional deviations. This implies that reaching efficient market outcomes requires no more than pairs of players making sure that they leave no gains from trade on the table,[21] providing formal foundations for extending the logic of the Coase Theorem to matching markets.

For the markets we consider, pairwise stable outcomes, or equivalently core outcomes, require that A is matched to C and B is matched to D for sure, while the combined payoffs of the strong players (A and D) must sum to weakly more than $x = 15, 25, 30$ for Game 15, Game 25 and Game 30, respectively. Thus, although the match is pinned down payoffs are not, and many different payoff profiles can be supported. In Appendix C we present the range of each player's payoffs that can be supported in a pairwise stable outcome.

Various theories have refined the set valued predictions provided by pairwise stability into point predictions. One alternative is to look at the mid-point of the supported payoffs (see, for example, Elliott (2015)).[22] A second alternative proposed by Rochford (1984), and independently by Kleinberg and Tardos (2008), extends Nash bargaining to matching markets. These symmetrically pairwise balanced (SPB) outcomes co-inside with several other cooperative solution concepts—specifically the nucleolus, kernel and pre-kernel. We develop these theoretical predictions for the markets we consider in Appendix C.[23]

3.5 Non-cooperative Theory

The experimental protocol we consider is standard and extends Rubinstein bargaining to accommodate many players. The corresponding game has an infinite-horizon with a common discount factor $\delta \in (0, 1)$. In round t there is a set of unmatched players who are active. One player is chosen uniformly at random to be a proposer. If the proposer is already matched, we move to round $t + 1$; otherwise the proposer can choose to propose a match or to do nothing (propose to oneself). To propose a match, the proposer must select an unmatched player and suggest a division of the surplus their match would generate. If a proposal is made, then the player who receives the proposal must either accept or reject it. If the proposal is accepted, then a match is formed and those two players, having reached agreement, leave the market. If the

[21] If the market outcome was not efficient, then the grand coalition would be able to form and implement the match that maximized total surplus, and then redistribute this surplus in a way that made everyone better off. Thus only efficient market outcomes are robust to coalitional deviations, and hence by Shapley and Shubik's result only efficient market outcomes are robust to pairwise deviations.

[22] This is the transferable utility equivalent to median stable matches in a non-transferable utility environment Schwarz and Yenmez (2011), which has received some experimental support in NTU matching experiments Echenique and Yariv (2013).

[23] The Shapley value makes unappealing predictions in matching markets, and so we do not consider it. For example, with one worker and two firms the Shapley value will typically require the firm which ends up unmatched to receive a transfer of surplus from the matched pair.

proposal is rejected, then both players remain unmatched and we move to round $t+1$. The game ends when there is no positive surplus between any two unmatched players.

Although there will often be multiple equilibria of this dynamic game, following the literature we focus on two criteria on which equilibrium selection can be based—simplicity and efficiency. Simplicity has led a large literature to study the **Markov perfect equilibria** (MPE) of related bargaining problems, including Rubinstein and Wolinsky (1985), Rubinstein and Wolinsky (1990), Gale (1987), Chatterjee and Sabourian (2000), Sabourian (2004), Gale and Sabourian (2006), Polanski and Winter (2010), Abreu and Manea (2012b), and Elliott and Nava (2015). In our context, the Markov perfect equilibria are perfect equilibria in which players choose strategies that depend only on which other players remain active in the market, rather than on the entire history of play.[24] While this prevents players from having to keep track of complicated histories of play, it also limits the ability of players to punish and reward each other.

When there are no efficient Markov Perfect equilibria, a natural question then arises as to whether more complicated strategies could obtain efficient outcomes. A second type of equilibria we will consider are **efficient perfect equilibria (PE)**. By design the markets we study require increasing complex strategies for an efficient perfect equilibrium. In Game 15 there is an efficient Markov perfect equilibrium. In Game 25 there is no efficient MPE, but there is an efficient PE that punishes deviations by reverting to the MPE. In game 30, there is no efficient MPE, or efficient PE that relies on Markov reversion, but there is an efficient perfect equilibrium that relies on more complicated strategies.

We now describe the limit MPE payoffs and provide some intuition for Games 15, 25 and 30. In all these games there is a unique MPE. In Game 15, all players always proposing efficiently is an MPE. When all players do so, it is as if they bargain bilaterally with their efficient partner and all players receive limit payoffs of 10. Given that offers of less than 10 will be rejected, it is unprofitable for A or D to deviate and instead offer to each other. Thus in Game 15 the efficient match is reached with probability 1.

In Game 25 it is no longer an equilibrium for only efficient offers to be made. If A and D never use their link they will get limit payoffs of 10 as before, but now they will have a profitable deviation to instead offer to each other when selected as the proposer. In equilibrium A and D mix between offering to each other and making efficient offers. Whenever A and D match with each other, players C and B get a payoff of 0. This reduces the amount players C and B are willing to accept when they do receive offers. In Game 25 the efficient match is obtained when C or B propose,

[24]The MPE are motivated in Maskin and Tirole (2001) and have been theoretically justified on complexity grounds as those selected when there is a second-order lexicographic preference for simple strategies (Sabourian, 2004).

but not always when A or D propose. Given the equilibrium probability with which A and D offer to each other, the probability the efficient match is reached is 0.72.

As the value of the diagonal link increases to 30 we reach a corner solution in which A and D can no longer push down the expected payoffs of C and B enough for them to be indifferent about whom to offer to. Hence A and D always offer to each other. Nevertheless, when selected as the proposer players C and B continue to make acceptable offers to A and D respectively, and we get the efficient match with probability 0.5.[25]

There are also equilibria in which non-Markovian strategies are played. In all our games an efficient perfect equilibrium exists, reflecting results in Abreu and Manea (2012a)[26].

There are two constraints that make finding an efficient perfect equilibrium hard. First, a player who makes an off-path offer cannot be punished if that offer is accepted (as the player exits). Second, in any efficient perfect equilibrium a subgame will be reached in which either just A and C are active or just B and D are active. In these subgames there is a unique subgame perfect equilibrium, and in this equilibrium both players' limit payoffs are 10. So once these subgames are entered there is no scope for rewards or punishments and all players receive relatively high payoffs.

In Game 25 there is an efficient PE supported by the threat of reverting to the MPE. We label these outcomes **efficient perfect equilibria with Markov reversion** (Eff. PE (i)). Interestingly, the threat of reverting to the MPE can only support on-path play that yields a unique vector of payoffs (see Section B in the Appendix). Constructing an efficient PE in Game 30 is more complicated. The threat of reverting to the MPE does not provide sufficient incentives to induce the strong players to offer to their efficient partners, but more complicated strategies can be used. In Section B in the Appendix we derive such strategies and show the range of payoffs they can support. We call these outcomes **efficient perfect equilibria with rewards and punishments** (Eff. PE (ii)). These strategies entail both rewards for not accepting offers that deviate from the prescribed play and punishments for deviating. Finally, we note that in Games 25 and 30 there does not exist an efficient perfect equilibrium in which the expected limit payoffs of players B and C sum to less than 10, which means that the average expected payoff of weak players must be at least 5.[27]

[25]For example, we can thus calculate the expected payoff of C as follows: With probability 0.5, A or D proposes and C gets 0; with probability 0.25 B proposes and reaches agreement with D leaving C to get 10 from bargaining bilaterally with A; and with probability 0.25, C proposes and gets 20 less the minimum offer A will accept.

[26]They show that by cleverly constructing punishments, an efficient perfect equilibrium always exists in markets where the gains from trade are either 1 or 0.

[27]The reason is that in any efficient perfect equilibrium either B or C will be left to bargain bilaterally with their efficient partner, thus, receiving a limit payoff of 10.

3.5.1 Comparison with Theoretical Predictions

Table 2: Predicted versus observed outcomes in the experienced games

	Result 1			Result 2			Result 3	
	(i)	*(ii)*	*(iii)*	*(i)*	*(ii)*	*(iii)*	*(i)*	*(ii)*
Coop.								
Core	✓	✗	✗	✓	✓	✗	✓	✗
SPB	✓	✗	✗	✓	✓	✗	✓	✗
Core Mid-Point	✓	✗	✗	✓	✓	✗	✓	✗
Non Coop.								
MPE	✓	✓	✓	✓	✓	✓	✓	✓
Eff. PE (i)	✓	✗	✗	✓	✓	✓	✓	✗
Eff. PE (ii)	✓	✗	✗	✓	✓	✓	✓	✗

Table 2 documents the qualitative performance of the different theories. All the theories do fairly well at describing how the payoffs of the strong and weak players vary across games and depend on whether they exit first or second (Result 2). However, the three theories perform differently in terms of their efficiency predictions (Result 1). While the efficient PE theories and cooperative theories predict no inefficiencies, the MPE predicts, as we observe, that there will be no inefficiencies in Game 15, but that inefficiencies then increase monotonically moving from Game 15 to Game 25, and then to Game 30. In terms of strategies (Result 3) the MPE also do a little better. The theories all correctly predict that players do not delay by making no offer, but the MPE is the only theory that predicts that strong players offer to each other with increasing frequency as we move from Game 15, to Game 25, to Game 30. Indeed, the MPE theory does relatively well quantitatively in this dimension.

In Figure 3 we present the empirical histogram of strong players' individual frequencies of proposing efficiently in each treatment. As the MPE predict that an agreement will always be reached in the first round we restrict attention to first round proposals.[28] The observed proposal frequencies, across the three treatments, fit the theory well.

Despite the agreement of our data with the above comparative static predictions, neither the MPE nor the efficient PE make quantitative predictions that match the data closely. The predictions of the different theories and the observed outcomes are presented in Table 3.[29]

[28] A figure that is very similar to Figure 3 is obtained when all rounds in which strong players have a choice to make are included. This figure is presented in the Supplementary Appendix, Section 1.4.

[29] While our non-cooperative theoretical predictions are derived for players as they become infinitely patient, i.e., for $\delta \to 1$, in the experiments we have implemented $\delta = 0.99$. This difference changes predictions in a very minimal way. None of our conclusions would change if we used $\delta = 0.99$ instead of letting $\delta \to 1$ to generate our theoretical predictions.

Figure 3: Frequency of efficient proposals by strong players in Experiment I, experienced games

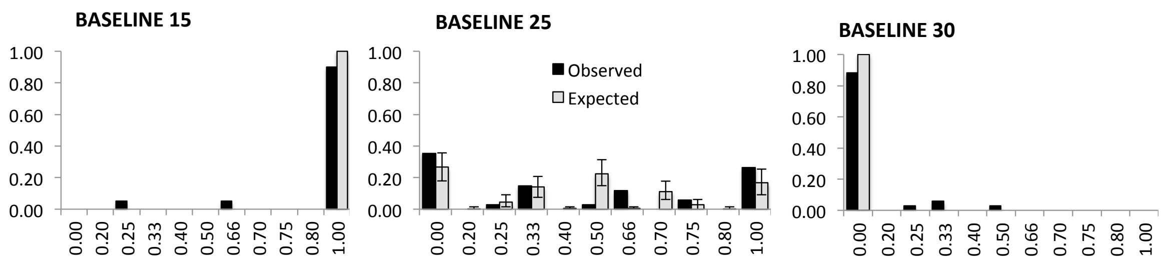

Notes: For each subject, we compute the frequency of proposing to her efficient partner in the first round of the last five repetitions conditional on this player being assigned a strong position. Black bars are the observed frequencies. Light grey bars take the distribution over the number of offers different players make from the data, and then simulate the frequencies with which these players make efficient offers assuming that they all play exactly the strategy prescribed by the MPE. We run 10,000 simulations and report the average frequencies as well as the range of frequencies between the 5th and 95th percentiles.

While the MPE theory is the only one of the three theories considered that predicts matching will be inefficient, the rate of mismatch we find in BASELINE 25 and BASELINE 30 treatments is substantially and statistically greater than is predicted by the MPE.

In neither Game 25 nor Game 30 does any theory predict payoffs for both weak and strong players within a 95% confidence interval of those observed. Moreover, despite the efficient perfect equilibria being able to support a range of payoffs (possibly even a larger range than the one we constructed if more complicated strategies are used) there does not exist an efficient perfect equilibrium that closely matches the observed payoffs in BASELINE 30 treatment. As discussed in Section 3.5, in any efficient perfect equilibrium the weak players must receive expected payoffs of at least 5, well outside the 95% confidence interval for the observed average payoffs of 2.4.

The average combined payoffs of players directly depends on the frequency with which an efficient match is reached. Given the frequent mismatches observed in BASELINE 25 and BASELINE 30 treatments, the MPE (which predicts mismatches) is the only theory that stands a chance of matching the observed payoffs of the weak and strong players. To strip away the efficiency dimension and focus on the division of the surplus between strong and weak players, we present the payoffs of players *conditional on reaching an efficient outcome* (last row of Table 3). For comparison, we also present the MPE-predicted payoffs of players who reach efficient outcomes (third row of Table 3). As evident from this table the MPE predicts players' payoffs within a 95% confidence interval in BASELINE 25, and it comes close to doing so in BASELINE 30. There also exist efficient PE strategies and core outcomes that generate consistent payoffs in both games .

Table 3: Predicted versus observed outcomes in Experiment I, experienced games

	Game 15			Game 25			Game 30		
	eff.	B (C)	A (D)	eff.	B (C)	A (D)	eff.	B (C)	A (D)
Coop.									
SPB	100%	8.3	11.7	100%	5	15	100%	3.3	16.7
Mid-Point	100%	10	10	100%	7.5	12.5	100%	5	15
Core	100%	[0,20]	[0,20]	100%	[0,15]	[5,20]	100%	[0,10]	[10,20]
Theory									
MPE all	100%	10	10	72%	6.45	11.45	50%	4.17	13.33
MPE | eff.		10	10		8.95	11.05		8.34	11.67
Eff. PE (i)	100%	10	10	100%	8.75	11.25	—	—	—
Eff. PE (ii)	100%	10	10	100%	$(7\frac{7}{9}, 9\frac{4}{9})$	$(10\frac{5}{9}, 12\frac{2}{9})$	100%	$(6\frac{1}{9}, 9\frac{4}{9})$	$(10\frac{5}{9}, 13\frac{8}{9})$
Data									
all	100%	10 (0.05)	10 (0.05)	51%	4.5 (0.36)	11.8 (0.13)	30%	2.4 (0.32)	14.2 (0.16)
| efficient		10 (0.05)	10 (0.05)		8.8 (0.17)	11.2 (0.17)		7.7 (0.37)	12.3 (0.37)

Notes: The last two rows report average payoffs of players by their network position, with the corresponding robust standard errors in the parenthesis where observations are clustered at the session level. The next-to-last row reports players' payoffs in all the final outcomes, while the last row focuses on the groups that reached an efficient outcome.

As we discuss in Section 6, players' strategies seem to vary over time, conditioning on the state of market, and so violate the history independence of the MPE. While clearly players are not playing exactly according to the MPE, overall the evidence suggests that the MPE organizes the data quite well. The data matches the MPE theory on many, albeit related, dimensions qualitatively, but differs from the quantitative predictions.

Our paper thus joins the emerging experimental literature that examines MPE in a variety of dynamic games (see Battaglini and Palfrey (2012), Battaglini et al. (2016), Salz and Vespa (2016), and Vespa (2016)). Similar to our results, in combination these papers show that comparative statics predictions implied by MPE organize experimental data quite well across a variety of dynamic games. Simplicity seems to be a useful guiding tool for equilibrium selection.

3.5.2 Takeaways from Experiment I

There are two key things that we takeaway from Experiment I. First, bargaining positions matter. Strong players, those with two links, do better than weak players, those with one link, and how much better depends on the value of their alternative match. Players' payoffs also depend on the market composition at the time they exit. Strong players do better if they reach agreement when both matches are still available to them. Second, even in our simple setting, markets very frequently fail to reach efficient outcomes.

This raises key questions about the possibility of the bargaining frictions we find in the lab, also being present in practice. In real world high skill markets, with limited search and informational frictions, will there still be a serious impediment to the efficient matches being reached in the form of bargaining frictions? If bargaining frictions do matter in such markets, might they also exacerbate problems in other markets more prone to search and informational frictions?

Addressing these questions is hard. The laboratory setting is removed from the real world in several ways. Nevertheless, better understanding what is driving the inefficiency we find in the lab can help inform us better about the likelihood of bargaining frictions having a substantial impact in real markets. To this end we advance two related explanations about what might be underlying the inefficiencies we find.

We know that whenever an inefficient match occurs, at least one worker and one firm have a profitable deviation they could exploit. So in our experiment there are pairs of players who are leaving surplus on the table. Moreover, there are simple heuristics in which profitable deviations are exploited and if followed lead the market to converge to a pairwise stable and efficient outcome (see, for example, Bayati et al. (2015)). Two related features of our experimental design could be causing problems that prevent this logic from being realized. First, the presence of rigid protocol governing who can interact with whom when might be preventing participants (in weak bargaining positions) who are about to be excluded, from making offers that would exploit a profitable pairwise deviation. For example, after a weak player's efficient partner has received an offer to match inefficiently this offer must be accepted or rejected before the weak player has the chance to put a counter proposal on the table. Second, agreements are reached sequentially and after a proposal has been accepted the players in question exit the market preventing them from exploiting profitable pairwise deviations that might remain. If instead players could renege on agreements, there would be no impediment to profitable pairwise deviations being exploited and efficient outcomes might be reached.[30]

Although our two explanations are closely related to each other, they have different implications for the conclusions that can be drawn from our experiment. If it is the experimental protocol that is generating the inefficiencies we find, then the external validity of our experiment would be quite limited. In the real world interactions are

[30]This explanation is also consistent with the the outside option principle from bargaining theory. Consider the inefficiencies we would expect to find were two players to bargain bilaterally, but with one of them having a long-lived outside option. In this case, the "outside option principle" predicts that the two players should reach agreement with probability 1 and, if the outside option is binding, the player with the binding outside option should receive a payoff equal to its value. Experimental support for precisely this is presented in Binmore et al. (1989). The reason the outside option principle does not apply to our experimental protocol is that it based on outside options always being available. When alternative matches can be lost, because others exit the market, they do not act like outside options. If players did not exit the market, and could renege on their current deals, then the same alternative matches would always be available.

not constrained by a protocol. If instead it is the sequential order in which agreements are reached, and the inability of people to renege on deals, external validity is less compromised. Jobs in labor markets are typically filled sequentially, while firms very rarely renege on agreements and it is unusual for people to do so too.[31]

To test these explanations and better understand what is driving inefficiencies in our first experiment, we run two new experiments. In the first experiment we allow the participants to interact freely with each other and do not impose an experimental protocol on interactions. In the second experiment we allow people to renege on existing agreements, for a small cost, while otherwise maintaining the protocol from the first experiment.

4 Experiment II: Unstructured Bargaining

Our first explanation for the inefficiencies we find in Experiment I is that it is driven by the highly structured nature of interactions as imposed by the specific experimental protocol used. While a variety of alternative bargaining protocols have been proposed in the literature,[32] in practice interactions in markets are much less structured than any of the candidate protocols from the theoretical literature. For example, it is typically endogenous whether an offer is responded to immediately with a counter-offer, or more generally who offers to whom when, and how long offers remain on the table before being withdrawn. Although there is not much hope for solving the dynamic game induced by such unstructured interactions, these games can be studied experimentally.

4.1 Experiment II: Experimental Design and Procedures

Experiment II consists of one treatment, UNSTRUCTURED 30, which uses market structure of Game 30 (see Figure 1). This treatment was conducted at the University of California, Irvine. A total of 88 subjects participated in three experimental sessions, which lasted less than one hour including the instruction period and quiz. Subjects earned on average $23.5 including a $15 show up fee.

Similarly to the BASELINE 30 treatment (Experiment I), in each session of UNSTRUCTURED 30 treatment, subjects play ten repetitions of Game 30 with random

[31] See, for example, Avery et al. (2001).

[32] In (Chatterjee et al., 1993) a player becomes the proposer after rejecting an offer, in (Abreu and Manea, 2012b) proposers cannot choose who they have the opportunity to make an offer to and in (Corominas-Bosch, 2004) everyone on the same side of the market simultaneously posts an offer. These games have different equilibrium outcomes and, beyond that, the protocol might affect equilibrium selection.

re-matching between games. Within each game there is no pre-determined structure of bargaining, instead subjects can make offers to any other player they might match to, withdraw offers they have previously proposed, and accept any currently standing offer they have received. Just like in the BASELINE 30 treatment, offers specify the player to whom it is made as well as the surplus split.[33] Once an offer is accepted, the match between these two players is formed and these players exit the market, i.e., have no more opportunities to move in the game. When a match is formed, all currently standing offers are voided, and bargaining starts afresh. Similarly to Experiment I, there are two ways in which the game can come to an end in Experiment II. First, the game ends if there are no new matches that can be formed between any two subjects who are not matched yet. Second, the game may end because of discounting, implemented as a random termination of the game: There is a 1% chance that the game ends at the end of each 30-second interval. When the game ends unmatched players receive a payoff of 0 while matched players earn payoffs according to their agreements. We refer the reader to Supplementary Appendix for the complete instructions used in these sessions as well as the screenshots and the quiz that subjects were asked to complete before the beginning of the experiment (Sections 7.1, 7.2, and 7.3).

4.2 Experiment II: Results

Our approach to the data analysis is the same as in the Experiment I (see Section 3.2.1). In particular, we focus on behavior in the experienced games (the last 5 repetitions of the game in each session) and perform most of the tests using regression analysis with standard errors clustered by session.

4.2.1 Market Outcomes

Our data shows that with unstructured bargaining markets continue to often match players inefficiently. Only 59% of the time in UNSTRUCTURED 30 treatment did the market clear efficiently. This is better than in BASELINE 30 treatment in Experiment I, but considerably less than all the time. With respect to players' payoffs, similarly to the BASELINE 30 treatment, we observe that the strong players obtain higher payoffs than the weak players ($p < 0.05$ in all comparisons). Moreover, as in BASELINE 30 treatment, in UNSTRUCTURED 30 treatment the order of exit also continues to

[33]Specifically, to make an offer a subject has to click on the ID letter of the subject she wants to receive the offer, and to type in the amount that she proposes to keep for herself; the remaining portion of the surplus the match would generate is allocated to the recipient. The offer then immediately appears in the column OFFERS YOU PROPOSED on the screen of the proposer and in the column OFFERS PROPOSED BY OTHERS on the screens of all the subjects still in the market. By clicking on a button the recipient can accept the offer at any time up until it is withdrawn.

matter. We present in Figure 4 the histogram of the final payoffs of the strong players in UNSTRUCTURED 30 treatment by their order of exit and conditional on an efficient outcome being reached. The grey bars in this figure depict payoffs of players that exited the market first and black bars present the payoffs of players that exited the market second.

Figure 4: Final payoffs of the strong players in efficient matches in Experiment II by market composition at the time of exit, experienced games

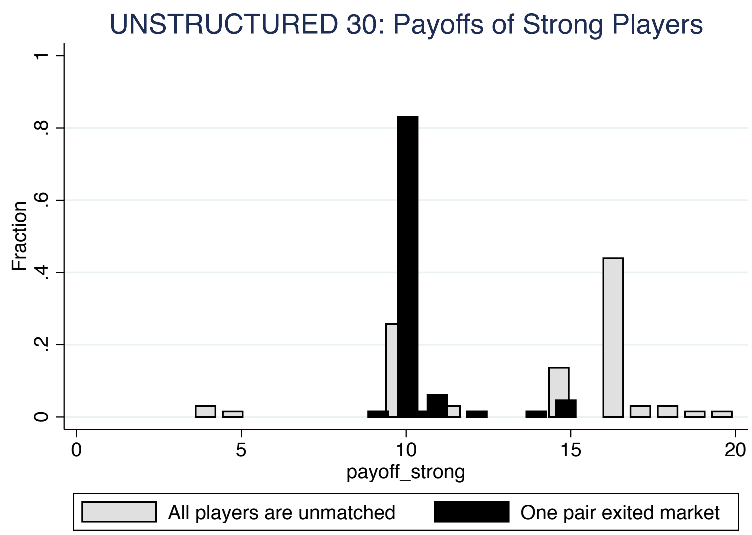

Notes: We consider only groups that reached efficient match.

The average payoff of strong players is 13.8 if they exited first, while it is 10.4 if they exited second, while the weak players earn on average 6.2 when exiting first compared to 9.6 when exiting second. Regression analysis confirms that these differences are statistically significant with $p < 0.05$ in all comparisons.

To sum up, when bargainers can act as they see fit in an unstructured manner, substantial inefficiencies persist and our data shows strong support for the market composition affecting the payoffs of players.

4.2.2 Takeaways from Experiment II

There are two main things we takeaway from Experiment II. First and foremost, despite allowing for unstructured bargaining we continue to find high rates of mismatch, albeit that they are a little lower than in Experiment I. Second, bargaining positions continue to matter. Players with two links do better than players with one link when both links are still available, but not once their alternative match has left the market.

There is a simple logic we expound in the introduction that can explain why it is hard to get efficient matching. Suppose there are two parties who match with each other in the efficient match, but that one of them is in a stronger bargaining position within the market than the other. If the efficient match is reached for certain, then the player in the weaker bargaining position can wait for others in the market to exit until she is no longer in a weak position. To prevent this, the player in the strong position has to sometimes matches inefficiently. In effect, the threat of matching with someone else is only credible if it is sometimes exercised. This logic is reflected in the results from Experiment II, as well as those from Experiment I. In both experiments players whose bargaining position might deteriorate sometimes match inefficiently, and doing so prevents their efficient partners from holding out for the opportunity to bargain bilaterally on an equal footing.

5 Experiment III: Reneging Experiment

We now test our hypothesis that if we limit the non-stationarities in bargaining positions, efficient outcomes will obtain. We test this hypothesis by running a third experiment in which we allow players to renege on existing agreements for a small cost.

5.1 Design and Procedures

The design and experimental procedures of Experiment III are very similar to those of the Experiment I. Experiment III consisted of 3 treatments: STAY 15, STAY 25 and STAY 30, with each treatment corresponding to one of the markets described in Figure 1.[34] Just like Experiment I, Experiment III was conducted in the same two locations: at the ESSL at University of California, Irvine and at the EBEL at University of California, Santa Barbara. Ten sessions were run, with a total of 156 subjects, recruited from a database of undergraduate students enrolled in these universities. The experiments lasted about two hours. Average earnings, including a $15 show up fee, were $23.7 with a standard deviation of $4.9.

In each session, subjects played ten repetitions of the same game, with random rematching between games. The main feature of Experiment III is the possibility of reneging on agreements formed in previous rounds. Recall that in Experiment I players have no opportunity of reneging, as those who reach agreements are forced to

[34]In this section, we briefly summarize the main features of Experiment III and refer the reader to the Supplementary Appendix, Sections 8.1, 8.2, and 8.3, in which we present instructions, the screenshots, and the quiz for this experiment.

exit the market permanently, which means they cannot make any further moves.[35] On the contrary, in Experiment III, players who reach agreements do not exit the market and can unilaterally break agreements they are part of at a small cost c.[36] In all three treatments of Experiment III, we used the same separation cost of $c = 10$ cents per broken agreement. Thus, a player who has formed a match remains active and can both propose new matches and accept new offers when proposed to her. If a currently matched player accepts a new offer, then she pays the separation cost for dissolving the previous match she was involved in and forms a new match in its place. If a currently matched player makes a new offer which is accepted by the responder, then the proposer pays the separation cost for breaking the match she was part of. The person who was part of an agreement that is broken by their partner in the current round does *not* pay the separation cost, but starts the next round unmatched. At the top of the screen subjects were reminded of the separation cost and of the number of times they have paid it in the current game.

All the remaining protocol details of Experiment III mirror those of Experiment I. In particular, there are two ways in which a game can end. First, there is a 1% chance that the game ends after each round, determined by a random draw of the computer. Second, the game ends if there was no positive surplus remaining between any pair of players who both made proposals in the last round.

5.2 Experiment III: Results

In this section we report several key comparisons of final outcomes and the strategies used by subjects in Experiment I and in Experiment III.

The premise of Experiment III is that the ability of players to renege on their current deals will make the bargaining environment stationary. To investigate the validity of this premise, we compare the monetary offers that unmatched strong players receive and accept, and the offers unmatched strong players make and have accepted, when (i) all players are unmatched, versus (ii) when the other strong player has already formed a tentative match with their efficient partner.

Table 4 shows that there is no evidence, in any of the three treatments STAY15, STAY25 or STAY30, that the offers received or made by an unmatched strong player, conditional on being accepted or not, depends on whether the other strong player was

[35] Specifically, when a player reaches an agreement with another player, the button responsible for submitting offers is disabled and the only active button on the screen of a matched player is the "Do Nothing" button, which she has to press in every round thereafter. We chose such a design in order to keep all the subjects engaged and focused on a game irrespective of the order in which they formed matches.

[36] In other words, the button responsible for submitting offers is never disabled no matter whether a player is matched with another player or not.

Table 4: Offers made to and received by unmatched strong players in Experiment III, experienced games

	Game 15			Game 25			Game 30		
	No match	Match	p-value	No match	Match	p-value	No match	Match	p-value
Offer made	10.57	10.43	0.52	12.98	12.30	0.09	15.54	15.41	0.71
And accepted	9.21	9.70	0.18	9.41	9.32	0.93	10.56	9.38	0.27
Offer received	9.71	9.31	0.30	11.77	11.67	0.79	13.58	13.14	0.46
And accepted	10.00	10.00	1.00	12.30	12.06	0.59	14.56	14.48	0.21

Notes: This table reports the efficient offers made by strong players (i.e., those to their efficient match), the accepted efficient offers made by strong players, the efficient offers received by strong players and the accepted efficient offers received by strong players. It does so for: (i) when all players are unmatched (no match); and (ii) when the other strong player is efficiently matched (match). The offers made by the strong players correspond to the amount the strong players proposed to keep for themselves. The offers received by strong players correspond to the payoffs the strong players get if these offers are accepted by them.

matched or not. This corroborates our hypothesis that, if reneging is possible, the bargaining positions of players remains stationary irrespective of whether the other players in the market are matched or not.

Our main hypothesis is in regard to a comparison of the efficiency levels observed in Experiment I with those in Experiment III. Figure 5 depicts the efficiency in each treatment in the experienced games, along with the 95% confidence intervals. As expected, in Game 15 we observe almost full efficiency regardless of the possibility of renegotiation. In the remaining two games, the possibility of reneging affects the final outcomes and significantly increases efficiency: In Game 25 efficiency increases from 51% to 82%, and in Game 30 it increases from 30% to 73%. While there is no statistical difference between the efficiency levels observed in the STAY 25 and STAY 30 treatments ($p > 0.10$), these payoffs are statistically less than 100% and so some inefficiency remains.

The increase in efficiency levels in Games 25 and 30 explains the differences in final payoffs of players between Experiment I and Experiment III reported in Table 5, where we report both average payoffs of players regardless of the final match and payoffs conditional on reaching an efficient match. Considering all final allocations, weak players earn on average higher payoffs in Experiment III than in Experiment I. This result is driven by the difference in efficiency levels, as weak players are less likely to remain unmatched when reneging is possible. However, conditional on reaching an efficient final match, strong players obtain higher shares, and weak players lower shares, in Experiment III compared with Experiment I. This is also intuitive. The possibility of reneging should make alternative matches more valuable, and it is the strong players who benefit from these alternatives.[37]

[37] In all regressions, we obtain p-values below the standard 5% level.

Figure 5: Efficiency of the final match in Experiments I and III, experienced games

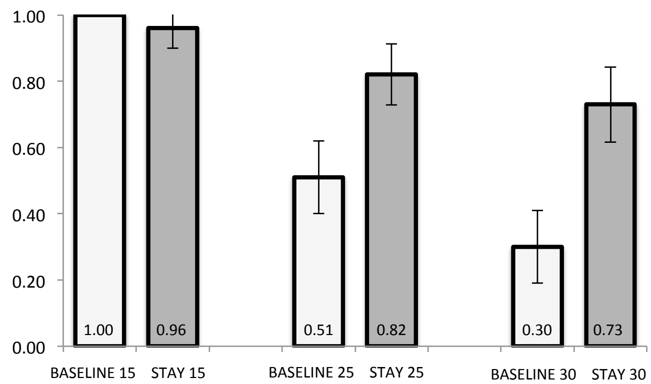

Notes: Average efficiency per treatment is reported, along with the 95% confidence interval, computed using robust standard errors, where errors are clustered at the session level.

Table 5: Payoffs of players by network position in Experiments I and III, experienced games

	Game 15		Game 25		Game 30	
	B (C)	A (D)	B (C)	A (D)	B (C)	A (D)
All Final Matches						
BASELINE treatment	10.0 (0.05)	10.0 (0.05)	4.5 (0.36)	11.8 (0.13)	2.4 (0.32)	14.2 (0.16)
STAY treatment	9.8 (0.16)	9.8 (0.16)	6.2 (0.37)	12.3 (0.16)	3.6 (0.22)	14.8 (0.15)
Efficient Final Matches						
BASELINE treatment	10.0 (0.05)	10.0 (0.05)	8.8 (0.17)	11.2 (0.17)	7.7 (0.37)	12.3 (0.37)
STAY treatment	10.0 (0.03)	10.0 (0.03)	7.5 (0.18)	12.3 (0.11)	4.9 (0.18)	15.0 (0.14)

Notes: We report average payoffs of players by their network positions, with the corresponding robust standard errors in the parentheses, where observations are clustered at the session level.

5.2.1 Takeaways from Experiment III

There are two main things we take away form Experiment III. First, bargaining positions continue to matter but unlike before there is no evidence that they change as agreements are reached. Our intention was to make the environment more stationary, and reassuringly the evidence is consistent with this. Second, and more importantly, with more stationary bargaining positions efficiency increases substantially, although non-trivial rates of mismatch remain.

6 Discussion

In this section we try to link our observed outcomes more closely to the theory and provide a coherent explanation for results across the experiments we run. This exercise is necessarily speculative. It lends itself to us trying to read too much into our results, while there are also likely to be alterative explanations we do not put forward. Indeed, we view our main contribution as demonstrating high levels of mismatch absent the frictions typically considered, and linking this to evolving bargaining positions. Nevertheless, it is interesting to consider the strategies of participants more carefully, and to hypothesize about how the theory might be adjusted to better fit our observations. We first explore the observed play of participants in Experiment I. This leads us to a relatively simple adjustment of the MPE theory. Building on this we look at (i) a possible explanation for the difference in the bargaining outcomes reached in Experiments I and II; and (ii) a possible explanation for the difference in the bargaining outcomes reached in Experiments I and III.

6.1 Players' Strategies in Experiment I

We are seeking to develop a new theory capable of matching the data quantitatively. To do this we will take the MPE as our starting point. While none of our subjects use strictly Markovian strategies,[38] in the experienced games, after learning has hopefully finished, the MPE do provide helpful guidance in organizing the experimental data. In Section 3.5.1 we found strong qualitative support, across many dimensions of the observed play in Experiment I, for the MPE. We thus begin by analyzing deviations from the MPE strategies.

The strategy of a player specifies a probability distribution over whom to make an offer to, details of such offers (amounts kept), and the minimum amount a player is willing to accept from others after every possible history of play. The restriction to Markovian strategies only allows players' strategies to depend on the state variable, which is the set of unmatched players. This greatly simplifies the strategy space, and, given the structures of the markets we consider, there are just two different states for us to analyze. First is the state in which all the players are unmatched and active, and the other is the one in which one of the efficient pairs has exited the market and the market consists of one remaining strong player and one remaining weak player.

When only two players remain active in the market, they both receive average

[38]Players' actions clearly do not only condition on the state variable (set of unmatched players), and other aspects of the history of play matter. This is, however, not very surprising given the plethora of experimental evidence that shows that learning and/or random errors are commonly observed in experiments on market interaction.

payoffs of 10, as the MPE predicts.[39] Moreover, as we discuss in Section 3.2.2, the frequency with which strong players offer inefficiently to other strong players in Experiment I tracks the MPE predictions across our three games (Game 15, Game 25 and Game 30). Thus, in the remainder of this section, we focus on the amounts players demand to keep when making offers as well as the amounts players accept and reject when all the players are active.

Table 6 depicts players' predicted ask amounts according to the MPE strategies as well as the observed ask amounts.

Table 6: Predicted and observed ask amounts in Experiment I when all players are unmatched

	BASELINE 15			BASELINE 25			BASELINE 30		
		data			data			data	
	MPE	1^{st} half	2^{nd} half	MPE	1^{st} half	2^{nd} half	MPE	1^{st} half	2^{nd} half
Strong to Strong	–	–	–	13.55	13.07	12.79	15.83	15.45	15.27
					(0.2)	(0.1)		(0.1)	(0.1)
Strong to Weak	10	11.28	10.54	13.55	12.64	13.26	–	–	–
		(0.2)	(0.2)		(0.2)	(0.4)			
Weak to Strong	10	10.99	10.62	8.55	9.30	9.06	6.67	9.53	7.88
		(0.2)	(0.2)		(0.4)	(0.4)		(0.4)	(0.2)

Notes: For the observed ask amounts, we report average ask amounts in the first round of the game broken down into the first and the second halves of the experiment, where the first half corresponds to the first 5 repetitions, and the second half to the experienced games. Robust standard errors are reported in the parenthesis, where observations are clustered by session. A − indicates that MPE predicts no offers of this kind.

On average, observed ask amounts deviate a bit from those predicted by MPE. The deviations are largest when weak players make offers to strong players. Weak players often make lower offers to strong players than predicted (so demand to keep more than predicted). If accepted, these offers would result in surplus being more equitably distributed between the strong and weak players. Indeed, we frequently observe weak players demanding perfectly equitable splits of 10 each in both BASELINE 25 and BASELINE 30 treatments. In BASELINE 25, for the experienced games when all players are active, 50% of weak subjects ask for an even split some times. In BASELINE 30, the corresponding percentage is 59%. Of the 50% in BASELINE 25 who sometimes propose equal splits, the average frequency with which an equal split is demanded is 60%. In BASELINE 30 it is 50%.

Remark 1. *In both BASELINE 25 and BASELINE 30 treatments, about 30% of the offers made by weak players demand a completely equitable split of* 10 *each, representing a substantial deviation from the MPE.*

A similar picture emerges from the analysis of offers that subjects accept and those they reject as responders. Consistent with MPE predictions, as the value of the

[39]Statistical analysis confirms this claim. In all three games we cannot reject the null that players' payoffs are 10 irrespectively of the network position with $p > 0.10$ in all cases.

diagonal link increases, strong responders tend to accept higher offers, while weak responders settle for lower shares. However, while the above qualitative difference between treatments is in line with the MPE, there are also some notable deviations. Strong players' acceptance strategies are roughly in line the MPE predictions. In BASELINE 25 they accept 84% of offers above $20 - 8.55 = 11.45$ (69 out of 82 cases) and reject 86% of offers below 11.45. In BASELINE 30 they accept 89% of offers above $20 - 6.67 = 13.33$ and reject 93% of offers below 13.33. At the same time, weak players' acceptance strategies are often at odds with the MPE. Theoretically, weak players should accept payoffs of $20 - 13.55 = 6.45$ in Game 25. In practice, they reject 50% of proposals from the strong players that offered them strictly less than equal split of 10, but more than 6.45.[40] In Game 30 weak players should never receive an offer and in line with this there are too few observations of weak players receiving offers for us to evaluate their acceptance strategies.

Can the deviations from MPE play that we have documented explain the increased inefficiency and inequality in the data relative to the MPE predictions? If weak players are relatively unlikely to reach agreement with strong players because they frequently demand too high a share of surplus, we will observe more inefficient agreements being reached more often than predicted. Thus the inefficient match will be implemented more often than predicted, and as weak players get a payoff of zero in this case, the payoffs of strong and weak players are also likely to be less equal than predicted. There is no statistical difference between the rate at which strong and weak players are selected to be the proposer, in BASELINE 25 and BASELINE 30 treatments, but the likelihood of reaching an agreement in the first round is significantly higher if the first mover is a strong player than if it is a weak player.[41]

6.2 Behavioral MPE in Experiment I

We have seen that weak players are frequently unwilling to reach inequitable agreements a player would be willing to reach in the MPE. In this section we propose a parsimonious adjustment to the MPE theory that incorporates this. While this adjustment to the theory will not be able to capture all of the rich behavior we observe, it incorporates an important driver of inefficiency missing from the MPE theory.

Taking inspiration from the large experimental literature analyzing two-person bargaining games (see Roth (1995) survey), and motivated by our analysis of players' strategies, we now discuss an adjustment of the MPE theory in which some players

[40] Offers below 6.45 are rejected 90% of the time with only 10 total observations of this type.

[41] In particular, in BASELINE 25 only 50% of offers made by weak players who were first movers were accepted, while this fraction is 70% for strong players who were selected to be first movers. Similarly, in BASELINE 30 only 53% of offers made by weak players who were first movers were accepted, while this fraction is 87% for strong players who were first movers.

always demand an equal shares of surplus with their efficient partner. Specifically, we assume that some players are rational while others are behavioral and only make offers that leave them with at least 10 (the equitable payoff in the efficient match), and only accept offers that leave them with at least 10.[42] We let each player's type be drawn independently, such that a given player is behavioral with probability x, and consider two similar adjustments to the MPE which are fully described in Section 3 of the Supplementary Appendix.

In the first adjustment, behavioral MPE I, we let players types be common knowledge. In Game 15 the behavioral players play equivalently to the rational players in the limit, so the adjustment makes no difference and the model continues to explain the data well. In Game 25 and Game 30 behavioral strong players also play in the same way as their rational counterparts, but as weak behavioral players play differently from weak rational players in these games, the types of players matters. When both weak players are rational, play is as before. When one weak player is behavioral and the other is rational, in equilibrium all players adjust their strategies relative to the rational benchmark in both Game 25 and Game 30. Finally, when both weak players are behavioral, the strong players only offer to each other while rejecting offers of 10 from their efficient partners and the inefficient match is reached for sure.

In the second adjustment, behavioral MPE II, players types are private information and we look for a Markov perfect Bayesian equilibrium. Off-path beliefs now need to be specified and we construct an equilibrium in which prior beliefs are retained, except in one case in which we consider doing so to be unnatural. We view our off-path beliefs are relatively intuitive, while they also limit the different beliefs we need to consider making the problem more tractable. It also helps that, like in behavioral MPE I, in Game 15 all behavioral players play equivalently to the rational players in the limit, while in Game 25 and Game 30 behavioral strong players also play in the same way as their rational counterparts. We do not attempt a complete characterization of all Markov perfect Bayesian equilibria.

To compare our predictions to the data we calibrate our model based on the proportion of behavioral players identified by our analysis of strategies in Section 6.1. Following Remark 1 we assume that 30% of our subjects are behavioral. Setting $x = 0.3$, we get the predictions shown in Table 7.[43]

The Behavioral MPE I predictions, in which players' types are known, are within the 95% confidence interval for the strong players' payoffs in Games 25 and 30, but not the weak players' payoffs or efficiency in either Game 25 or Game 30. Nevertheless, introducing behavioral types in this way does tend to move predictions closer to the data. The Behavioral MPE II predictions, in which players' types are not known,

[42] We would expect similar theoretical predictions if instead behavioral players demanded 50% of the surplus available in the match in question, whatever that match is.

[43] We report predictions for a wide range of values of x in Section 3 of Supplementary Appendix.

Table 7: Predicted versus observed outcomes in Experiment I, experienced games

	Game 15			Game 25			Game 30		
	eff.	B (C)	A (D)	eff.	B (C)	A (D)	eff.	B (C)	A (D)
MPE	100%	10	10	72%	6.45	11.45	50%	4.17	13.33
Behavioral MPE I	100%	10	10	65%	5.4	12.0	53%	3.4	14.2
Behavioral MPE II	100%	10	10	44%	3.8	12.0	39%	3.0	13.9
Data	100%	10	10	51%	4.5	11.8	30%	2.4	14.2
		(0.05)	(0.05)	(0.06)	(0.36)	(0.13)	(0.06)	(0.32)	(0.16)

Notes: The behavioral MPE predictions are reported for when 30% of players are behavioral. The last row report average payoffs of players by their network position, with the corresponding robust standard errors in the parenthesis where observations are clustered at the session level.

move predictions closer still. These predictions are (just) within the 95% confidence interval for everything: efficiency, the weak players' payoffs, the strong players' payoffs in both Game 25 and Game 30. Overall we view the behavioral adjustment to take the quantitative predictions of the theory considerably closer to the data.

Interestingly, and in contrast to the ultimatum game where evidence for the presence of these players has also been documented, the introduction of behavioral types results in more unequal outcomes as well as more inefficient outcomes. Strong players respond to the demands of behavioral weak players by more frequently excluding them from agreements, yielding more mismatch and inequality. Norms that do well limiting inequality in the ultimatum game exacerbate inequality when alternative matches are available.

6.3 Structured Vs. Unstructured Bargaining

Our analysis of Experiment I shows that the presence of behavioral players can explain the higher-than-predicted rates of mismatch, in comparison to the MPE. In this section, we build on this theory to provide a (speculative) explanation for rates of mismatch we observe in Experiment II. A first step for doing so is to examine the results of this experiment in more detail.

Table 8 documents the observed outcomes from Experiment II (unstructured bargaining) and reproduces the theoretical predictions for ease of comparison. While the cooperative theories predict that markets always reach the efficient match, our data shows that markets often fail to do so. Despite being based on a bargaining protocol disconnected from the unstructured interactions, qualitatively the MPE continue to do well. There is inefficient matching as predicted, the payoffs of the weak players are less than than those of the strong players, and as shown in Figure 4, the payoffs of strong players are higher when they exit first.[44] Surprisingly, the MPE now also do well quantitatively. Observed efficiency levels of 59% are quite close and not statistically different at the 5% level from the efficiency levels predicted by MPE for Game 30, which is precisely 50%. The payoffs predicted by the MPE, for both weak and strong players, are also within 95% confidence intervals of those observed.

Table 8: Outcomes in Experiment II, experienced games

	Game 30		
	eff.	B (C)	A (D)
Cooperative approach			
SPB	100%	3.3	16.7
Mid-Point	100%	5	15
Core	100%	[0,10]	[10,20]
Non-Cooperative approach			
MPE	50%	4.17	13.33
MPE | eff.		8.34	11.67
Data			
all	59%	4.7 (0.57)	13.2 (0.23)
| efficient		7.9 (0.36)	12.1 (0.36)

Notes: The last two rows report average payoffs of players by their network position, with the corresponding robust standard errors in the parenthesis where observations are clustered at the session level. The next-to-last row reports players' payoffs in all the final outcomes, while the last row focuses on the groups that reached an efficient outcome.

We have seen that the presence of behavioral players can explain the results from Experiment I. By demanding an equal shares of surplus despite being in weak bargaining positions, these players are left unmatched more frequently than they would be in the MPE. Moreover, the presence of behavioral players depresses the average payoff of weak players, while leaving the payoffs of strong players unaffected—so if their interest is in equitable outcomes their actions are self-defeating. It seems possible that these players might learn to adjust their behavior.

However, we hypothesize that it is hard to learn this in the structured environment. The theory predicts that conditional on matching, weak behavioral players receive higher payoffs than their rational counterparts do. Thus, when these players successfully match, as occurs with positive probability, it might reenforce their behavior. Moreover, rational weak players are frequently left unmatched, it is just that behavioral weak players are left unmatched more often. In comparison, in the unstructured bargaining experiment, if a strong player holds an offer from the other strong player, the weak players know that the receiving strong player can accept this offer at any time (up until this offer is withdrawn), and that doing so will result in

[44] We do not study the prediction from the first experiment related to who strong players offer to, because in the protocol-free environment such choice are defunct—strong players can offer to both possible matches at the same time.

the weak players getting a payoff of zero. We conjecture that this could lead weak behavioral players to adjust their demands and make their efficient partners acceptable offers. Moreover, it is now less costly for a strong player to make an offer to their efficient partner. They can always make an offer the other strong player at the same time. Nothing prevents a strong player from having alternative offers on the table at once (if one offer is accepted the other is immediately and automatically withdrawn). Both mechanisms can help efficient matches be reached by behavioral players who are willing to adjust their demands.

With both structured bargaining (BASELINE 30) and unstructured bargaining (UNSTRUCTURED 30) weak players frequently propose an equal split of surplus. Recall that in BASELINE 30 treatment all players make proposals each round, and then one is selected uniformly at random to be implemented. Considering all the proposals, including those not implemented, in the experienced games (last 5 repetitions), 52% of the groups had at least one weak player propose an equal split when markets were complete (such that no players had exited). With unstructured bargaining, in the experienced games, 57% of groups had at least one weak player propose an equal split when markets were complete. However, groups involving weak players who demand equal splits are much more likely to reach the efficient match with unstructured bargaining than structured bargaining. In the experienced games of BASELINE 30, only 6% of those groups in which a weak player demanded an equal split ultimately reached the efficient match. With unstructured bargaining 61% of those groups involving a weak player who proposed an equal split reached the efficient match. Moreover, weak players receive many more offers in the unstructured experiment. With structured bargaining (BASELINE 30), more than 85% of strong players never offer to a weak player, while with unstructured bargaining (UNSTRUCTURED 30), the average frequency with which strong players offer to a weak player is 67%. This is all consistent with the richer interactions in the unstructured experiment permitting behavioral weak players to learn that they will likely be unmatched if they insist on an equal split of surplus, and adjusting their demands accordingly.

6.4 Inefficiency: Experiment I versus Experiment III

In comparison to Experiment I, efficiency increases substantially in Experiment III, although non-trivial rates of mismatch remain (see Figure 5). The increase in efficiency is driven by two mechanisms. First, strong players are significantly more likely to make offers to their efficient partners in comparison to Experiment I (Appendix D). Second, more efficient matches are also reached through profitable deviations to mismatches being exploited after the inefficient match has first been implemented. In both STAY 25 and STAY 30 treatments there are more than 30% of groups in which strong players first formed an inefficient match between themselves, only to then realize profitable deviations and rematch with their efficient partners. While this behavior is more consistent with the pairwise stability theory than the MPE theory, mismatch is not eliminated altogether (in contrast to the pairwise stability theory).

We consider players' strategies to better understand why mismatch remains. For the STAY 25 treatment suppose one efficient match has been agreed in which the strong player receives an amount x. Then, by the logic of pairwise stability, the unmatched weak player should be willing to accept a payoff of $y = 20 - (25 - x)$. Any agreement in which the weak player gets more than this would leave the two strong players with a profitable deviation. We look for evidence that weak players reject offers greater than equal to y. We find that such offers are rejected in the experienced games in 9 out of 25 cases (36% of the time). We do the same exercise for STAY 30. In this case, we'd expect the unmatched weak player to accept a payoff of $y = 20 - (30 - x)$. We find that, offers to unmatched weak players greater than or equal to y are rejected in the experienced games on 5 out of 17 occasions (36% of the time).[45]

Having found evidence that some weak players demand higher payoffs than the theory predicts they should, we explore whether this can account for the inefficiency we observe. To do this we introduce behavioral players. As when considering Experiment I, we suppose that a fraction x of players are behavioral and demand a payoff of at least 10. We then adjust the pairwise stability theory accordingly. We exclude stable matches in which behavioral players receive a payoff less than 10 and include unstable matches in which the only profitable deviations require a behavioral player to receive a payoff of less than 10. We set the proportion of behavioral player to be $x \in \{0.3, 0.4\}$, and present the resulting theoretical predictions in Table 9 where Behavioral Core I corresponds to $x = 0.3$ while Behavioral Core II is for $x = 0.4$.

While there do appear to be some players in our data who demand exactly 10,[46] there are also others who seem willing to accept lower payoff than 10, but not payoffs as low as they should be willing to accept. We do not claim that the simple adjustment to the pairwise stability theory that we make perfectly captures the rich behaviour of our subjects. However, it does incorporate into our theory a key missing feature—that weak players sometimes demand more than they should. As we show below, this can explain why profitable deviations are not exploited and why some inefficiency persists.

The derivation of the predictions in Table 9 are in Section 5 of the Supplementary

[45]In STAY 25, when unmatched weak players make offers, they offer their efficient partner strictly less than $25 - x$ in 194 out of 247 cases (79% of the time). In STAY 30, when unmatched weak players make offers, they offer their efficient partner strictly less than $30 - x$ in 159 out of 237 cases (67% of the time).

[46]In STAY 25 treatment, when market is complete, weak players propose to keep 10 out of 20 in 37% of the cases in the experienced games, while in STAY 30 treatment they do so in 15% of the cases.

Table 9: Predicted versus observed outcomes in Experiment III, experienced games

	Game 15			Game 25			Game 30		
	eff.	B (C)	A (D)	eff.	B (C)	A (D)	eff.	B (C)	A (D)
Core	100%	[0,20]	[0,20]	100%	[0,15]	[5,20]	100%	[0,10]	[10,20]
Behavioral Core I	100%	[0,20]	[0,20]	91%	[0,15]	[5,20]	[91%,49%]	[0,10]	[10,20]
Behavioral Core II	100%	[0,20]	[0,20]	84%	[0,15]	[5,20]	[84%,36%]	[0,10]	[10,20]
Data	100%	9.8	9.8	82%	6.2	12.3	73%	3.6	14.8
		(0.16)	(0.16)		(0.37)	(0.16)		(0.22)	(0.15)

Notes: The behavioral MPE predictions are reported for when $x = 30\%$ and $x = 40\%$ of players are behavioral. The last row report average payoffs of players by their network position, with the corresponding robust standard errors in the parenthesis underneath average payoffs where observations are clustered at the session level.

Appendix. In Game 25 inefficient outcomes are reached if and only if both players are behavioral. When both weak players are behavioral the strong players can never reach an agreement with their efficient partners that would not leave them with a profitable deviation to instead match to each other. However, when only one of the weak players is behavioral, the pairwise stable match is efficient. The behavioral weak player then receives a payoff of at least 10, his efficient partner receives the remaining surplus from the match, the other strong player receives at least 15, and the other weak players receives at most 5.

In Game 30, when there are two behavioral players the only stable match is again the inefficient one. However, when there is one behavioral player, both the efficient and inefficient match can be supported in a pairwise stable outcome. In the former case the weak behavioral player and his partner can receive 10 each, while the other strong player receives 20 and his efficient partner, the other weak player, receives 0. In the later case, the strong players can be matched together, with the strong player who's efficient partner is the behavioral weak player receiving 10 and the other strong player receiving 20. As both these outcomes are possible, a wide range of efficiency levels is consistent with the theory for this game.

7 Conclusions

Market clearing is a fundamental question in economics. It is important to get the "right" people into the "right" jobs, especially in the high-skill labor markets in which mismatches can be very costly in term of efficiency. In this paper we remove standard frictions and study mismatch. We turn off search frictions by letting players be patient, remove information problems by giving everyone symmetric information about match surpluses, remove coordination problems by considering very simple markets and give norms a good chance of yielding the efficient outcome by considering markets in which a perfectly equitable and efficient outcome is feasible. We still find persistent and extensive mismatch across our three experiments. Inefficiencies are highest when interactions are constrained by an experimental protocol and participants are not permitted to renege on agreements they reach. Removing the bargaining protocol, thereby permitting much richer endogenous interactions, inefficiency improves a little, but remains substantial. Instead permitting agents to renege on agreements, making the environment more stationary, efficiency improves more, but non-trivial inefficiencies persist. As bargaining positions seem non-stationarity in many high-skill labor markets, our results are consistent with the bargaining frictions we document contributing towards mismatch in actual markets.

To further investigate our findings we seek to adjust existing theory so that it can fit the data quantitatively, as well as qualitatively. Although this exercise is certainly speculative, and it is likely that other theories could be constructed to explain this or other aspects of our data, our conjectured explanation presents a consistent picture across the experiments we run and is motivated by a careful examination of the strategies we observe.

In our first experiment, participants' strategies are suggestive of a particular behavioral bias. Similarly to previous experimental work on the ultimatum game, some players in weak positions demand a more equitable split of surplus than theory suggests they should. However, unlike in the ultimatum game, in our setting an adjusted equilibrium that accounts for the presence of these behavioral players, results in less equal outcomes. Instead of being offered more, behavioral players are left unmatched more often, and market outcomes are both less efficient and less equitable than when more behavioral players are present.

When bargaining is protocol-free, as in Experiment II, our data is consistent with the richer interactions that are now permitted allowing players in weak positions to learn that demanding equitable outcomes is a bad idea—insofar as it reduces their own expected payoffs as well as leading to less equitable outcomes overall—and then adjusting their demands.

When, instead, reneging is permitted, as in Experiment III, the strategies of players seems more in line with the pairwise stability theory, than the MPE. The same behavioral adjustment to this theory, in which some weak players demand a more equitable split of surplus that the theory predicts they should, generates predictions consistent with our data.

Taking this interpretation of our results at face value, there are two sources of inefficiency in our first experiment. The first comes through what we refer to as bargaining frictions. Fundamental non-stationarities in the strength of players' positions leads rational players to match inefficiently. The second comes through behavioral players who are unwilling to accept inequitable splits of surplus. When these players find themselves in weak bargaining positions their demands lead to them being excluded from the market. Permitting reneging helps remove the bargaining frictions, but not behavioral frictions. Permitting richer interactions through protocol-free bargaining helps remove the behavioral frictions, but not the bargaining frictions.

References

Abreu, D. and M. Manea (2012a): "Bargaining and efficiency in networks," *Journal of Economic Theory*, 147.

——— (2012b): "Markov equilibria in a model of bargaining in networks," *Games and Economic Behavior*, 75, 1–16.

Avery, C., C. Jolls, R. A. Posner, and A. E. Roth (2001): "The market for federal judicial law clerks," *The University of Chicago Law Review*, 68, 793–902.

Battaglini, M., S. Nunnari, and T. Palfrey (2016): "The Dynamic Free Rider Problem: A Laboratory Study," *American Economic Journal: Microeconomics*.

Battaglini, M. and T. Palfrey (2012): "The Dynamics of Redistributive Politics," *Economic Theory*, 49, 739–777.

Bayati, M., C. Borgs, J. Chayes, Y. Kanoria, and A. Montanari (2015): "Bargaining dynamics in exchange networks," *Journal of Economic Theory*, 156, 417–454.

Bienenstock, E. J. and P. Bonacich (1993): "Game-theory models for exchange networks: experimental results," *Sociological Perspectives*, 36, 117–135.

Binmore, K., A. Shaked, and J. Sutton (1989): "An outside option experiment," *The Quarterly Journal of Economics*, 104, 753–770.

Calvo-Armengol, A. and M. O. Jackson (2004): "The effects of social networks on employment and inequality," *The American Economic Review*, 94, 426–454.

Charness, G., M. Corominas-Bosch, and G. Fréchette (2007): "Bargaining and network structure: An experiment," *Journal of Economic Theory*, 136, 28–65.

Chatterjee, K., B. Dutta, D. Ray, and K. Sengupta (1993): "A noncooperative theory of coalitional bargaining," *The Review of Economic Studies*, 60, 463–477.

Chatterjee, K. and H. Sabourian (2000): "Multiperson bargaining and strategic complexity," *Econometrica*, 1491–1509.

Choi, S., A. Galeotti, and S. Goyal (2014): "Trading in networks: Theory and experiments," *Cambridge-INET Working Paper*, 8.

Choi, S., E. Gallo, and S. Kariv (2016): "Networks in the laboratory," *Oxford Handbook on the Economics of Networks, forthcoming.*

Cook, K. S. and R. M. Emerson (1978): "Power, equity and commitment in exchange networks," *American Sociological Review*, 721–739.

Cook, K. S., R. M. Emerson, M. R. Gillmore, and T. Yamagishi (1983): "The distribution of power in exchange networks: Theory and experimental results," *American Journal of Sociology*, 275–305.

Corominas-Bosch, M. (2004): "Bargaining in a network of buyers and sellers," *Journal of Economic Theory*, 115, 35–77.

Echenique, F. and L. Yariv (2013): "An experimental study of decentralized matching," *mimeo.*

Eeckhout, J. and P. Kircher (2011): "Identifying sorting—in theory," *The Review of Economic Studies*, rdq034.

Elliott, M. (2015): "Inefficiencies in Networked Markets," *American Economic Journal: Microeconomics*, 7, 43–82.

Elliott, M. and F. Nava (2015): "Decentralized bargaining: Efficiency and the core," *mimeo.*

Featherstone, C. and E. Mayefsky (2010): "Stability and deferred acceptance: Strategic behavior in two-sided matching," *mimeo.*

Gale, D. (1987): "Limit theorems for markets with sequential bargaining," *Journal of Economic Theory*, 43.

Gale, D. and S. Kariv (2009): "Trading in networks: A normal form game experiment," *American Economic Journal: Microeconomics*, 114–132.

Gale, D. and H. Sabourian (2006): "Markov equilibria in dynamic matching and bargaining games," *Games and Economic Behavior*, 54, 336–352.

Hall, R. E. and A. B. Krueger (2012): "Evidence on the incidence of wage posting, wage bargaining, and on-the-job search," *American Economic Journal: Macroeconomics*, 56–67.

Haruvy, E. and M. U. Ünver (2007): "Equilibrium selection and the role of information in repeated matching markets," *Economics Letters*, 94, 284–289.

Jackson, M. O. (2010): *Social and economic networks*, Princeton university press.

Kagel, J., R. A. E. (2000): "The dynamics of reorganization in matching markets: A laboratory experiment motivated by natural experiments," *Quarterly Journal of Economics*, 115, 201–235.

Kalai, E. and W. Stanford (1988): "Finite rationality and interpersonal complexity in repeated games," *Econometrica*, 397–410.

Kleinberg, J. and É. Tardos (2008): "Balanced outcomes in social exchange networks," in *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, ACM, 295–304.

Kranton, R. and D. Minehart (2001): "A theory of buyer–seller networks," *American Economic Review*, 91, 485–508.

Maskin, E. and J. Tirole (2001): "Markov perfect equilibrium: I. Observable actions," *Journal of Economic Theory*, 100, 191–219.

Molis, E., V. R. (2010): "Experimental results on the roommate problem," *Journal of Economic Theory*, 145, 2218–2240.

Nalbantian, H. R. and A. Schotter (1995): "Matching and efficiency in the baseball free-agent system: An experimental examination," *Journal of Labor Economics*, 13, 1–31.

Niederle, M. and A. E. Roth (2009): "Market culture: How rules governing exploding offers affect market performance," *American Economic Journal: Microeconomics*, 2, 199–219.

Pais, Joana, A. P. and R. F. Veszteg (2011): "Decentralized matching markets: A laboratory experiment," *mimeo.*

Palfrey, T. (2016): "Experiments in political economy," in *The Handbook of Experimental Economics*, vol. II, forthcoming.

Polanski, A. and E. Winter (2010): "Endogenous two-sided markets with repeated transactions," *BE Journal of Theoretical Economics*, 10.

Rochford, S. C. (1984): "Symmetrically pairwise-bargained allocations in assignment markets," *Journal of Economic Theory*, 34, 262–281.

Rogerson, R. and R. Shimer (2011): "Search in macroeconomic models of the labor market," *Handbook of Labor Economics*, 4, 619–700.

Rogerson, R., R. Shimer, and R. Wright (2005): "Search-theoretic models of the labor market: A survey," *Journal of Economic Literature*, 43, 959–988.

Roth, A. (1987): "Equity and coalition bargaining in experimental three-person games," in *Laboratory Experimentation in Economics: Six Points of View*, ed. by A. Roth, Cambridge University Press.

——— (1995): "Bargaining experiments," in *The Handbook of Experimental Economics*, ed. by J. Kagel and A. Roth, Princeton Univ. Press.

Rubinstein, A. (1982): "Perfect equilibrium in a bargaining model," *Econometrica*, 50, 97–109.

Rubinstein, A. and A. Wolinsky (1985): "Equilibrium in a market with sequential bargaining," *Econometrica*, 53, 1133–1150.

——— (1990): "Decentralized trading, strategic behaviour and the Walrasian outcome," *The Review of Economic Studies*, 63–78.

Sabourian, H. (2004): "Bargaining and markets: Complexity and the competitive outcome," *Journal of Economic Theory*, 116, 189–228.

Salz, T. and E. Vespa (2016): "Estimating Dynamic Games of Oligopolistic Competition: An Experimental Investigation," *working paper*.

Schwarz, M. and M. B. Yenmez (2011): "Median stable matching for markets with wages," *Journal of Economic Theory*, 146, 619–637.

Shapley, L. and M. Shubik (1972): "The assignment game," *International Journal of Game Theory*, 1, 111–130.

Skvoretz, J. and D. Willer (1993): "Exclusion and power: A test of four theories of power in exchange networks," *American Sociological Review*, 801–818.

Vespa, E. (2016): "An Experimental Investigation of Strategies In Dynamic Games," *working paper*.

A Experiment I: Markov Perfect Equilibria

In this section we derive the MPE predictions summarized in Section 3.3. A more formal and general derivation is provided in Elliott and Nava (2015). We start with Game 25. Let $W(\delta)$ be the continuation value of players in the subgames where they are bargaining bilaterally with their efficient partners. By Rubinstein (1982) there is a unique perfect equilibrium in these subgames and $\lim_{\delta\to 1} W(\delta) = 10$. Letting V_i be the continuation value of player i when no one has yet been matched, we look for a symmetric solution in which $V_A = V_D := V_S$ is the continuation value for both the strong players and $V_B = V_C := V_W$ is the continuation value for both the weak players. We guess and verify that in Game 25 there is an equilibrium in which the strong players mix between offering to each other and offering to their efficient partners. Letting q be the probability that that the strong players, A and D, offer inefficiently to each other if either is selected as the proposer, we then have the following system of equations.

$$\begin{aligned}
V_S &= \frac{1}{4}\left(20 - \delta V_W + \delta(1+q)V_S + (2-q)\delta W(\delta)\right) \\
V_W &= \frac{1}{4}\left(20 - \delta V_S + (1-q)\delta V_W + (2-q)\delta W(\delta)\right) \\
20 - \delta V_W &= 25 - \delta V_S,
\end{aligned}$$

The first two equations state the continuation values of the strong and weak players as determined by the possible transitions in states that can occur and the payoffs associated with these transitions. The last equation is an indifference condition that must be satisfied for the strong players to strictly mix who they offer to.

Solving this system of equations and taking limits,

$$\begin{aligned}
q &\to \frac{16 - \sqrt{160}}{6} = 0.56 \\
V_S &\to 11.45 \\
V_W &\to 6.45
\end{aligned}$$

Given these continuation values, it is easily verified that no players have a profitable deviation. As players are always offered their continuation values, acceptance is optimal, and as the strong players mix, by construction they are indifferent between offering to each other and offering to their efficient partners. Finally, delaying is unprofitable. In the limit, by deviating and delaying a weak player receives an expected payoff of $6.45 < 20 - 11.45$, while a strong player receives an expected payoff of $11.45 < 20 - 6.45 = 25 - 11.45$.

For Game 30, players A and D strictly prefer offering to each other. The system of equations is then

$$\begin{aligned}
V_S &= \frac{1}{4}\left(30 - \delta V_S + 2\delta V_S + \delta W(\delta)\right) \\
V_W &= \frac{1}{4}\left(20 - \delta V_S + \delta W(\delta)\right)
\end{aligned}$$

Solving this system of equations and taking limits, we get

$$
\begin{aligned}
V_S &\to \frac{40}{3} = 13.33 \\
V_W &\to \frac{25}{6} = 4.17
\end{aligned}
$$

It is again easily verified that this is an equilibrium. For example, were a strong player to deviate and make an offer to a weak player, the lowest acceptable offer they could make would leave the strong player with a payoff of $20 - 4.17 < 30 - 13.33$.

B Experiment I: Efficient Perfect Equilibria

To construct an efficient perfect equilibrium, we need to create the right system of rewards and punishments for players to play efficiently. One measure of the complexity of strategies is the extent to which the players' prescribed actions vary with the history of play (see, for example, Kalai and Stanford (1988)). In this sense Markovian stratagies are particularly simple, as they depend only on the history through the state—in this case the set of active players. In order to incentivize the players to play efficiently, more complicated strategies are necessary to create the right system of rewards and punishments.

We start by considering a particularly simple class of efficient perfect equilibria, those where (only) reversion to the Markov perfect equilibrium is used as a punishment. Thus equilibrium play depends only on the state and whether there has been a deviation. It doesn't matter who deviated or when. In Game 25 there is a perfect equilibrium that can be supported by reversion to the MPE. Let the expected MPE payoff of player i be $V_i^M(\delta)$, and as before let $W(\delta)$ be the payoff of a player in the unique perfect equilibrium of the subgame where all other players except her efficient partner has exited.

We construct an efficient perfect equilibrium in which, on path, player i offers her efficient partner $\mu^*(i)$ a payoff δV_i^M and the offer is accepted. Any deviation from this play is punished by moving to the Markov perfect equilibrium. Thus, by construction, player i best responds by accepting the offer. Indeed, given that deviations are supported by reversion to the MPE, player i must offer her efficient partner exactly $\delta V_{\mu^*(i)}^M$. Anything less would be rejected by $\mu^*(i)$. If the strategy prescribed i offering anything more than $\delta V_{\mu^*(i)}^M$ to $\mu^*(i)$, then i would have a profitable deviation to offer a little less and, knowing that because i has deviated play will revert to the MPE strategies thereafter, $\mu^*(i)$ would accept.

It is easily verified that $\delta V_i^M < 20 - \delta V_{\mu^*(i)}^M$ for all players and thus all players prefer making the prescribed offers to delaying. The final deviation to check is that the strong players cannot do better by offering to each other. In Game 25 this requires that $25 - \delta V_S^M \leq 20 - \delta V_W^M$. As a strong player offering to another strong player constitutes a deviation, thereafter the MPE will be played. Hence each strong player will just be willing to accept an offer from the other strong player that leaves her with a payoff of δV_S^M. In the MPE the strong players mix between making offers to each other and offering to their efficient partner, and so are indifferent between these alternatives implying that $25-\delta V_S^M = 20-\delta V_W^M$. Thus the above inequality is satisfied and the strong players do not have a profitable deviation. However, this is not the case for Game 30. In Game 30 the strong players do not have a profitable deviation if $30 - \delta V_S^M \leq 20 - \delta V_W^M$ (where these MPE continuation values are for Game 30 and not for Game 25 as before). As in the MPE of Game 30 the strong players strictly prefer offering to each other than offering efficiently, $30 - \delta V_D^M > 20 - \delta V_C^M$. Thus in Game 30 Markov reversion does not provide sufficient incentives for the strong players to offer efficiently and there is no efficient MPE with Markov revision for Game 30.

In the efficient PE with MPE reversion for Game 25, the limit payoffs of the players are

$$
\begin{aligned}
V_A &= \frac{1}{4}\big(20 - \delta V_C^M + \delta V_A^M + \delta 2W(\delta)\big) \to 11.25 \\
V_B &= \frac{1}{4}\big(20 - \delta V_D^M + \delta V_B^M + \delta 2W(\delta)\big) \to 8.75 \\
V_C &= \frac{1}{4}\big(20 - \delta V_A^M + \delta V_C^M + \delta 2W(\delta)\big) \to 8.75 \\
V_D &= \frac{1}{4}\big(20 - \delta V_B^M + \delta V_D^M + \delta 2W(\delta)\big) \to 11.25
\end{aligned}
$$

To further illustrate that there is no such efficient PE for Game 30, and letting V_i^M now refer to the expected MPE payoff of player i in game 30 (as opposed to Game 25 above), the limit payoffs of the players would be

$$
\begin{aligned}
V_A &= \frac{1}{4}\big(20 - \delta V_C^M + \delta V_A^M + \delta 2W(\delta)\big) \to 12.29 \\
V_B &= \frac{1}{4}\big(20 - \delta V_D^M + \delta V_B^M + \delta 2W(\delta)\big) \to 7.71 \\
V_C &= \frac{1}{4}\big(20 - \delta V_A^M + \delta V_C^M + \delta 2W(\delta)\big) \to 7.71 \\
V_D &= \frac{1}{4}\big(20 - \delta V_B^M + \delta V_D^M + \delta 2W(\delta)\big) \to 12.29
\end{aligned}
$$

But then if selected as the proposer, A can either stick with the prescribed strategy

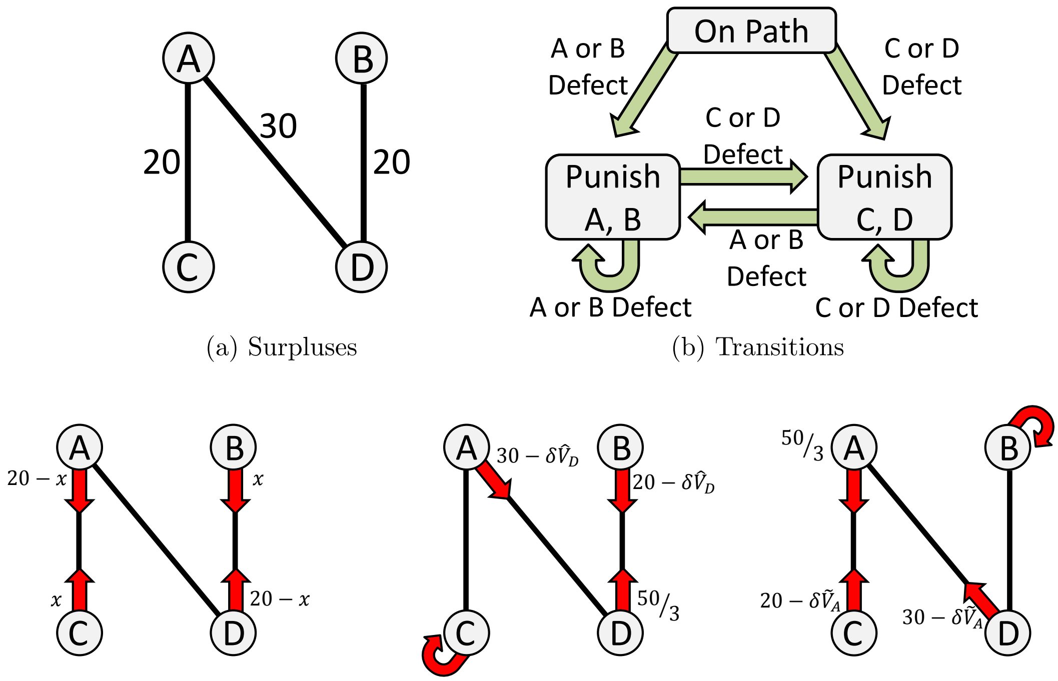

(a) Surpluses (b) Transitions

(c) On path (d) Punish A,B (e) Punish C,D

Figure 6: Constructing an efficient perfect equilibrium for Game 30. Panel (b) shows the transitions between states when players deviate from the prescribed play, while panels (c)–(e) show how players play in each state. Red arrows indicate whom a player offers to if selected as the proposer; and the numbers next to the arrows indicate the payoffs that the offering players will keep.

that offers C a payoff δV_C^M, leaving A with a limit payoff of 15.83, or deviate and offer D a payoff δV_D^M, which D would accept, leaving A with a limit payoff of 16.66.

To find a limit MPE for Game 30 we need to consider more complicated strategies, in which players are both rewarded for rejecting off-path offers and punished for making off-path offers. The rewards are important because we can punish an off-path offer only if it is rejected. On-path, we look for an equilibrium in which players B and C make efficient acceptable offers that leave them with a payoff of x and accept offers of x from their efficient partners. This is shown in panel (c) of Figure 6. On-path, after the first efficient pair of players exit the market the remaining efficient pair bargain bilaterally with each other. In such subgames there is a unique perfect equilibrium (Rubinstein, 1982) and the remaining active players receive payoffs $W(\delta)$ that converge to 10. Thus, in any efficient equilibrium the last weak player to reach agreement receives a limit payoff of 10. In order to get these weaker players to accept and make offers that give them a payoff of $x < 10$ we need to punish them if they deviate.

We construct off-path punishments that are credible and create the appropriate incentives for players to remain on path. This is achieved by defining two different punishment states, prescribing play in each of these states and a rule for transitioning between them in a way that creates the appropriate incentives. These transitions are such that they occur only if someone deviates from their prescribed strategy, in which case the person who initiated the deviation is punished by moving to the state that punishes her. Importantly, these transitions also reward all the players to whom the punished player is linked. These transitions are illustrated in panel (b) of Figure 6.

To show that the punishments are credible, suppose we are in the Punish A, B state. If everyone plays as prescribed we remain in this state and the payoffs of the players are given by the following value functions:

$$
\begin{aligned}
\widehat{V}_A &= \frac{1}{4}\left(30 - \delta\widehat{V}_D + 2\delta W(\delta) + \delta\widehat{V}_A\right) = \frac{30 - \delta\widehat{V}_D + 2\delta W(\delta)}{4-\delta} \to 11\frac{1}{9} \\
\widehat{V}_B &= \frac{1}{4}\left(20 - \delta\widehat{V}_D + 3\frac{1}{3} + \delta\widehat{V}_B\right) = \frac{20 - \delta\widehat{V}_D + 3\frac{1}{3}}{4-\delta} \to 2\frac{2}{9} \\
\widehat{V}_C &= \frac{1}{4}\left(\delta\widehat{V}_C + 2W(\delta)\right) = \frac{2W(\delta)}{4-\delta} \to 6\frac{2}{3} \\
\widehat{V}_D &= \frac{1}{4}\left(16\frac{2}{3} + 3\delta\widehat{V}_D\right) = \frac{16\frac{2}{3}}{4-3\delta} \to 16\frac{2}{3}
\end{aligned}
$$

By symmetry, the punish C, D state value functions of the players are

$$
\begin{aligned}
\widetilde{V}_A &= \frac{16\frac{2}{3}}{4-3\delta} \to 16\frac{2}{3} \\
\widetilde{V}_B &= \frac{2W(\delta)}{4-\delta} \to 6\frac{2}{3} \\
\widetilde{V}_C &= \frac{20 - \delta\widehat{V}_D + 3\frac{1}{3}}{4-\delta} \to 2\frac{2}{9} \\
\widetilde{V}_D &= \frac{30 - \delta\widehat{V}_D + 2W(\delta)}{4-\delta} \to 11\frac{1}{9}
\end{aligned}
$$

Consider now the deviations available to the players in the punish A, B state. First, suppose that A deviates and offers D less than $\delta\widehat{V}_D$. By rejecting the offer, D ensures that we remain in the same state and that he will receive, in expectation, $\delta\widehat{V}_D$. Alternatively, A may delay, in which case A receives $\delta\widehat{V}_A < 30 - \delta\widehat{V}_D$. Finally, A could offer to C. C would accept anything greater than $\delta\widehat{V}_C$ and reject anything less, because we would remain in the punish A,B state. Thus A must offer C a limit payoff of $6\frac{2}{3}$, leaving A with $13\frac{1}{3} \leq 30 - 16\frac{2}{3} = 13\frac{1}{3}$.[47] The alternative deviations available to B are to offer D less than $\delta\widehat{V}_D$, which D would reject, leaving B with $2\frac{1}{3} < 3\frac{1}{3}$, or to delay, which would also leave B with $2\frac{1}{3} < 3\frac{1}{3}$. The only deviation available to C is to make an offer to A. As rejecting C's offer will result in a switch of states, A would accept only a limit payoff which is weakly greater than $16\frac{2}{3}$, leaving C with $3\frac{1}{3} < 6\frac{2}{3}$. Finally, D could deviate. As a deviation by D would result in a switch of states if rejected, for an off-path offer to be accepted D must offer B at least $6\frac{2}{3}$ in the limit, or A at least $16\frac{2}{3}$ in the limit. Both deviations are thus unprofitable. Finally, as delay would also result in a switch of states, that alternative is unprofitable for D as well. This covers all the possible deviations from the punish A, B state. By symmetry, there are no profitable deviations from the punish C, D state.

For these punishments to be effective, in the on-path state C and B must be required to accept only offers, in the limit, of weakly more than $2\frac{2}{9}$ in the limit, or to make offers that leave them with at least $2\frac{2}{9}$ in the limit. Similarly, A and D must be required to accept only offers of weakly more than $11\frac{1}{9}$ in the limit, or to make offers that leave them with at least $11\frac{1}{9}$ in the limit. Thus for any $x \in (2\frac{2}{9}, 8\frac{8}{9})$ there exists an efficient perfect equilibrium. This places bounds on the offers that can be supported when all the players are active. In the subgame reached once an efficient pair has exited, the remaining players get limit payoffs of 10. Thus the weak players will have limit expected payoffs in the range $(6\frac{1}{9}, 9\frac{4}{9})$, while the strong players will have limit expected payoffs in the range $(10\frac{5}{9}, 13\frac{8}{9})$.

The construction we use to find an efficient perfect equilibrium for Game 30 also works for Game 25. In that case, in the punish A, B state we would need to make A offer D no more than $11\frac{2}{3}$, leaving A with $13\frac{1}{3}$, so that A does not want to deviate and instead offer to C. As A offers D his continuation value, this implies that $V_D = 11\frac{2}{3}$, which means that B also offers $11\frac{2}{3}$ to D and that D offers $8\frac{1}{3}$ to B. This gives a limit payoff to B of $V_B = 5\frac{5}{9}$. As before, C's limit payoff is $V_D = 6\frac{2}{3}$. Finally, A's limit payoff is $11\frac{1}{9}$. Given these strategies and limit payoffs, it can be verified that all the incentive constraints are satisfied. Thus, there is an efficient perfect equilibrium for any $x \in (5\frac{5}{9}, 8\frac{8}{9})$. As limit payoffs in the subgame are again 10, weak players have limit expected payoffs in the range $(7\frac{7}{9}, 9\frac{4}{9})$, while strong players have limit payoffs in the range $(10\frac{5}{9}, 12\frac{2}{9})$.

[47] This inequality would becomes strict, while the others remain satisfied, were the offer D makes to B to be increased slightly. We make this incentive constraint tight to make the punishment as harsh as possible so that we can find the full range of payoffs that can be supported.

C Cooperative Theories

With unstructured bargaining, efficiency enhancing norms could emerge to the mutual benefit of everyone. The participants have much more freedom to reach implicit agreements that exploit inefficiencies for their mutual gain. This logic is present in several solution concepts from cooperative game theory, and with unstructured bargaining we might expect their predictions to be born out. To guide our experiment we consider two such theories: (i) an extension of Nash bargaining where the disagreement points are defined endogenously by the market; and (ii) the mid-point of the core.[48]

The first solution concept we consider is **Symmetric Pairwise Bargained** (SPB) outcomes, first developed by Rochford (1984) and independently discovered by Kleinberg and Tardos (2008). This approach extends Nash bargaining to networks. A player's disagreement payoff is the surplus they could obtain by just enticing someone else to match with them. Given these disagreement payoffs, two players reaching agreement each get their disagreement payoff plus an equal share of the remaining surplus. Of course, the disagreement payoff for a given player depends on the agreements others reach, and so the solution boils down to finding a fixed point of a large system of equations. We derive predictions for Game 15, 25 and 30 in the Supplementary Appendix Section 4.1. These predictions co-inside with the predictions of three other cooperative solution concepts, the kernel, pre-kernel and nucleolus.

Our next cooperative theory is an alternative refinement of the core. The core comprises the set of bargaining outcomes that allocate the surplus generate by matching such that there is no coalition of players that could find a match among themselves, and then allocate that surplus in a way that makes them all better off. Shapley and Shubik (1972) show that there are no profitable coalitional deviations if and only if there are no profitable pairwise deviations, which just requires that for all workers firm pairs that the sum of their payoffs is weakly greater than the surplus they could generate by matching with each other. This implies that the match implemented must be efficient. It also follows by results from Shapley and Shubik (1972) that there is a **mid-point of the core** (mid-point) in which two players reaching an agreement receive their worst possible core payoff for sure, and then split the remaining surplus equally.[49] In a non-transferable utility (NTU) environment, Echenique and Yariv (2013) find experimentally that the median stable match is reached. For the markets

[48]We do not consider the Shapley value or related concepts. The Shapley value makes predictions that can be infeasible in our setting when each matched pair of players is required to split between themselves the surplus they generate. For example, if there is one worker and two firms and all matches generate a positive surplus the Shapley value requires that all players receive strictly positive payoffs.

[49]These payoffs are special cases of Corominas-Bosch (2004) and Elliott (2015) while Kranton and Minehart (2001) use a related selection.

we consider, the mid-point of the core corresponds to the transferable utility (TU) median stable matching (Schwarz and Yenmez, 2011). These predictions are derived in the Supplementary Appendix Section 4.2 and reported below in Table 10.

Table 10: Theoretical predictions about final matches

	Game 15			Game 25			Game 30		
	eff.	B (C)	A (D)	eff.	B (C)	A (D)	eff.	B (C)	A (D)
SPB	100%	8.3	11.7	100%	5	15	100%	3.3	16.7
Core	100%	[0,20]	[0,20]	100%	[0,15]	[5,20]	100%	[0,10]	[10,20]
Core Mid-Point	100%	10	10	100%	7.5	12.5	100%	5	15

D Further Analysis of Experiment III

In this Appendix we ask whether the shift in efficiency levels documented in Figure 5 is driven by players rematching when gains from trade are left on the table, or if it comes from a change in the strategies used by players as a response to a change in the environment. For each game, Figure 7 compares the CDFs of individual frequencies, in the STAY and BASELINE treatments, of proposing efficiently by strong players in the experienced games.[50] As is evident from this figure, except for Game 15, in which the vast majority of subjects always propose efficiently, strong players propose efficiently with higher frequencies when there is a possibility of renegotiation.[51] Regression analysis confirms these results: $p < 0.01$ in both the BASELINE 25 vs. STAY 25 regression and the BASELINE 30 vs. STAY 30 regression, while $p > 0.10$ in BASELINE 15 vs. STAY 15 regression.

To further examine how more efficient matches are reached we classify the efficient matches in STAY 25 and STAY 30 treatments according to whether or not they went through a mismatch (a match between two strong players) before reaching the efficient matches. Focusing on the experienced games, we observe that the majority of groups in both the STAY 25 and STAY 30 treatments reached an efficient final outcome right away, without ever mismatching; there are 68% and 63% of groups like this in the

[50]In Figure 7 we focus on the first-round behavior of strong players in each repetition in the experienced games. Including all rounds in which all players were unmatched generates a figure very similar to Figure 7 (see the Supplementary Appendix, Section 1.4). Note that data presented in Figure 6 regarding strategies used by the strong players in the BASELINE treatment is a different representation of the same data presented in Figure 3 (CDF instead of histogram).

[51]For example, in the BASELINE 25 treatment more than 50% of the players in strong positions proposed efficiently less than one third of the time, while in the STAY 25 treatment less than 50% of the players in strong positions proposed efficiently less than 60% of the time.

Figure 7: CDFs of frequency of efficient proposals by strong players in Experiment III, experienced games

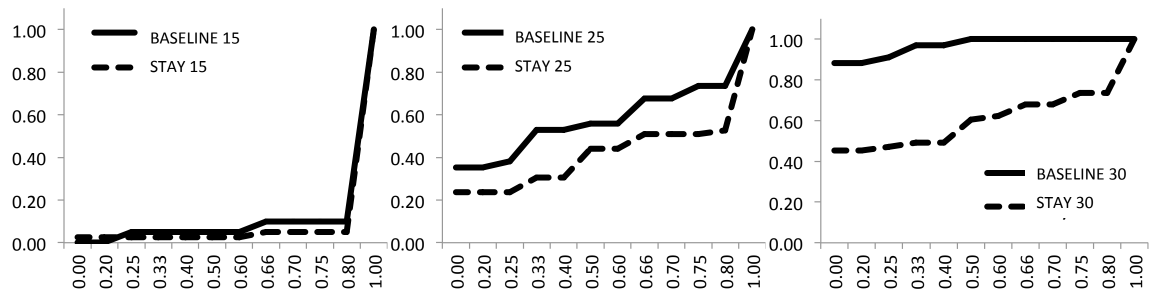

Notes: We present the cumulative distribution functions summarizing individual frequencies of proposing efficiently in the first round of the last five repetitions in each session when a subject performed a role of strong player. The horizontal axes indicate the likelihood of proposing efficiently, while the vertical axes indicate the values of the CDFs.

STAY 25 and STAY 30 treatments, respectively. Correspondingly, in both treatments there are more than 30% of groups in which strong players first formed an inefficient match between themselves, only to then realize profitable deviations and rematch with their efficient partners.[52] We conclude that while much of the efficiency gain can be attributed to players making efficient proposals more often, more efficient matches are also sometimes reached through profitable deviations to mismatches being exploited.

[52]Only one group in STAY 25 treatment and three groups in the STAY 30 treatment went through forming the inefficient match twice.