Commitment and (In)Efficiency: a Bargaining Experiment*

Marina Agranov[†] Matt Elliott[‡]

January 29, 2016

Abstract

We conduct an experimental investigation of decentralized bargaining over the terms of trade in matching markets. We study if/when efficient matches will be reached and what terms of trade will be agreed. Multiple theories guide our analysis and we test their predictions against the outcomes of our experiments. We find that inefficiencies are extensive and they are driven by the endogenous evolution of bargaining positions as agreements are reached and players exit. When we allow subjects to renege on existing agreements efficiency significantly improves.

1 Introduction

A fundamental question in economics is whether the right people end up in the right jobs. It is crucial for the understanding of labor markets and the productivity of the economy. In many labor markets multiple workers and multiple employers simultaneously bargain with each other to determine who will be employed by whom and on what terms. As agreements are reached the composition of the market, those workers and firms who remain actively searching for each other, changes. As this market context evolves, so can the bargaining positions of the remaining workers and

*We thank Federico Enchenique, Edo Gallo, Ben Gillen, Ben Golub, Francesco Nava, Emanuel Vespa, Leeat Yariv, and participants in seminars at Columbia, NYU, UCSD, Georgetown, Berkeley HAAS 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`.

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

firms. When can such labor markets be expected to clear efficiently and match the right workers to the right firms? How does the composition of the market affect the terms agreed?

Theoretical predictions in relation to the aforementioned questions are mixed, and these are hard questions to investigate empirically. To know the extent of mismatch an econometrician must estimate the counterfactual productivities of matching different people to different jobs. But unobservable characteristics can generate any counterfactual productivities and rationalize any given match as efficient. To overcome this problem, we take an experimental approach. We use an array of experiments to study how the structure of the market (who can generate what surplus with whom) impacts on efficiency and payoffs. In the lab, we can control the entire set of match surpluses. We can also track individuals' bargaining patterns in full and, consequently, assess the outcomes that they yield.

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 wage bargaining as opposed to non-negotiable wages being posted (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.[1]

We seek to better understand whether the evolving nature of markets can cause workers and firms to match inefficiently and how payoffs are affected. To avoid countervailing effects, we do so in the simplest possible setting. We assume matching is one-to-one (so, in the context of labor markets, firms have a single vacancy to fill), consider markets with two workers and two firms, abstract from informational problems by assuming these possible match surpluses are common knowledge, and assume there is transferable utility. In this context, our experimental subjects engage in decentralized bargaining to determine who will be matched to who and on what terms. An efficient match is defined as one that maximizes the combined payoffs of the players.

We consider three different markets, defined by the surpluses different worker-firm pairs can achieve, and these differ only in the surplus available to one worker-firm pair. The markets we consider are shown in Table 1, where we set $x = 15$ for the first market, $x = 25$ for the second market and $x = 30$ for the third market. In all three cases it is efficient for worker A to match with firm C and for worker B to match with firm D. By varying x we vary the strength of the bargaining positions of the "strong players" (worker A and firm D) relative to the "weak players" (worker B and firm C).

[1] http://www.heritage.org

Table 1: Surpluses

		Workers	
		A	B
Firms	C	20	0
	D	x	20

A standard and parsimonious way to model market clearing in matching markets is to utilize cooperative game theory and abstract from the dynamics. For example, Rochford (1984) and Kleinberg and Tardos (2008) show that the Nash bargaining solution can be extended to market contexts by defining players' disagreement payoffs as the highest payoff they could obtain by making someone else an offer that would leave her no worse off. This implies that matches will be efficient. Moreover, even weak alternatives, such as the option of worker A matching to firm D to generate a surplus of 15 in our first market, will improve these players' bargained outcomes. These symmetric pairwise bargained (SPB) outcomes always yields matches and payoffs in which no pair of players has a profitable deviation.

An alternative cooperative approach is to directly look at the set of outcomes that are pairwise stable and make a different selection. While in all such outcomes matches are efficient, a range of payoffs can be supported. There is a pairwise stable outcome in which all workers simultaneously receive their lowest possible pairwise stable payoff and all firms simultaneously receive their maximum possible stable outcome, known as the firm-optimal payoffs (Shapley and Shubik (1972)). Worker-optimal payoffs are defined analogously and also exist. Moreover, the efficient match with any convex combination of these payoffs is also pairwise stable. For any such convex combination, in our first market with $x = 15$ it is predicted that there is no difference between the strong and weak players, in contrast to the SPB outcomes.

While, starting from Becker (1973), the cooperative approach has proved popular,[2] when the dynamics of the market are modeled complicated strategies can be required to reach efficient equilibrium outcomes (Gale and Sabourian (2006), Abreu and Manea (2012b), Abreu and Manea (2012a), Elliott and Nava (2015)).[3] The types of strategies necessary to reach efficient outcomes depend on the market structure. In some cases simple (stationary / Markovian) strategies suffice, in others reversion to simple strategies provides sufficient incentives, yet in some cases complex rewards and punishments must be constructed to incentivize the players to match efficiently.

[2]For example, Cole et al. (2001a); Cole et al. (2001b); Elliott (2015) and Elliott (2014) use this framework to study investment incentives.

[3](Shimer and Smith, 2000) extends Becker (1973) to incorporate dynamic matching in the presence of search frictions in large markets. In contrast, we focus on thin markets where an individual agreement can change the market composition in a meaningful way but players' are sufficiently patient to limit search frictions.

Our experiment is designed to differentiate between these three cases. In our first market, with $x = 15$, there exists an efficient Markov perfect equilibrium (MPE). In the second market, with $x = 25$, there is no efficient MPE but there is an efficient perfect equilibrium in which reversion to the MPE is used to provide incentives. In the third markets, with $x = 30$, there is no MPE and no perfect equilibrium with reversion to the MPE, but there is a perfect equilibrium in which more complicated rewards and punishments are used.[4]

The headline results of our experiments show that (1) inefficiencies in the form of mismatches are common whenever alternative matches provide binding outside options (when $x = 25$ and $x = 30$), and (2) efficiency declines as one moves from the first market ($x = 15$) to the second market ($x = 25$) to the third market ($x = 30$). These trends are consistent with MPE predictions but not with any of the other theories we consider, although mismatch in our data occurs more frequently than MPE predicts. On the payoffs dimension, observed payoffs of players across the three markets change in a way consistent with most of the theories, while none of the theories generates point predictions that consistently match the data across the markets.[5] Beyond the final outcomes, the observed market dynamics and the evolution of bargaining power before final outcomes are reached is largely consistent with the MPE predictions within and across the three markets.

Motivated by excessive mismatches observed in our first experiment, we set out to further explore the mechanism driving these inefficiencies. We hypothesize that, in line with the MPE logic, it is the prospect of bargaining positions changing over time that drives inefficiencies. Specifically, we conjecture that were we to allow players to renege on existing matches with a vanishingly small cost, then there would always exist an efficient MPE. By allowing agents to rematch at a small cost, alternative matches become more stationary. With alternatives remaining available, the environment is much closer to one in which the market provides fixed outside options. Thus we conjecture that, in line with the "outside option principle," agents should be able to agree to terms of trade disciplined by their alternative matches without having to exercise them.[6] In other words, there should always exist an efficient MPE. While there are significant impediments to proving this conjecture, we take a first step to-

[4]The three cases thus differ in how many different "punishment states" are needed to incentivize the agents to match efficiently. For the first market none are needed, for the second market one is needed and for the third market two are needed. Kalai and Stanford (1988), for example, use this as a measure of strategic complexity.

[5]The alternative of A matching to D in the market with $x = 15$ does not help A and D achieve higher payoffs than B and C, providing further grounds for rejecting the SPB prediction.

[6]This principle asserts that only binding outside options should matter and, when they do bind, the worker and firm will still reach agreement just with the split of surplus determined by the binding outside option. In other words, the threat of taking the outside option is enough to discipline payoffs without ever being exercised. The genesis of this idea is in Ståhl (1972), and there is a nice exposition closer to our setting in Sutton (1986).

wards evaluating it experimentally. To this end we conduct a second experiment, in which we explore the same three markets but allow players to rematch at a small cost. We find that efficiency increases significantly in these new treatments. This is driven by two forces. First, as conjectured, agents are more likely to make an offer to their efficient partner than to someone else. Second, in contrast with the conjecture, agents sometimes get to the efficient match by first mismatching and then exploiting profitable deviations.[7]

1.1 Related Literature

We focus in this section on the related experimental literature. We discuss the theoretical literature in Section 3.

There is a large experimental literature on bargaining.[8] 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 document that the two players reach agreement and responders receive a payoff equal to their binding outside option, providing support for the "outside option principle."

Our paper naturally relates to the experimental literature on decentralized two-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 branch include Enchenique and Yariv (2013) and Pais and Veszteg (2011). Enchenique 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

[7]We view this second channel as being more in line with the cooperative theories we consider than with our conjecture. These theories rule out outcomes in which a pair of agents has a profitable deviation, implicitly assuming such deviations would be realized.

[8]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.

final outcomes.[9]

The main difference between our paper and those discussed above is that we are interested in 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 consider matching markets in which agents have private information about their payoffs and analyze several procedures for matching. These matching procedures range from the free-agency system, which replicates 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.[10] 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 value of a homogeneous and indivisible good. 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.[11]

[9]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.

[10]The free-agency system is implemented using a clever design in which market participants have a fixed amount of time to strike deals with each other using a voice-to-voice protocol. The participants are seated in different offices and can bargain with each other via phones. As agreements are reached, they are publicly announced and the remaining players continue bargaining until they reach agreements or the time is over.

[11]See also Gale and Kariv (2009) and Choi et al. (2014) for a study of trading in networks with

Finally, there is an experimental literature in sociology that studies how network structures confer power. Two foundational papers in this literature are Cook and Emerson (1978) and Cook et al. (1983). A typical experimental design in this literature has several features different from us, including that negotiations are unstructured and that all trades are equally valuable. The focuses is on identifying strong and weak network positions. Perhaps the paper closest to ours is Skvoretz and Willer (1993), who test four sociological theories, including some from cooperative game theory that we consider. Their environment features unstructured trades that are equally valuable, and their focus is wider than our matching markets. Players cannot be partitioned into buyers and sellers, some players can trade multiple units, and each pair has a maximum capacity of trades they can make. As far as we are aware, the sociology literature does not investigate the interaction between changing market composition and efficiency.

2 Environment

2.1 Setup

We set out to test whether 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.

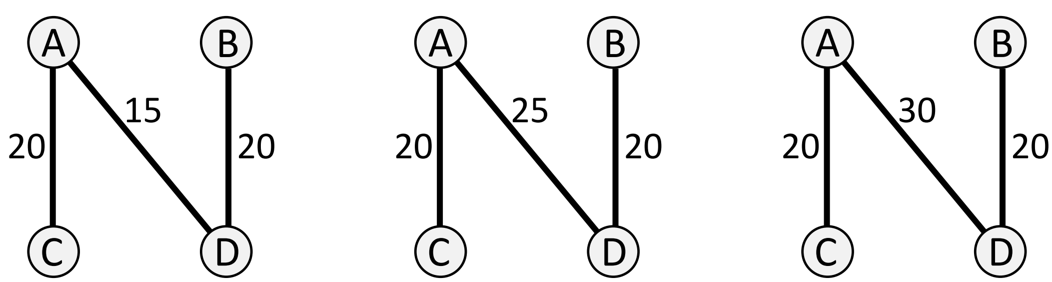

(a) Game 15 (b) Game 25 (c) Game 30

Figure 1: The three networks considered in this study. We refer to players A and D as the strong players and players B and C as the weak players.

Figure 1 presents three different market structures (Game 15, Game 25, and Game

intermediaries, implemented through a simultaneous bid-ask protocol and posted prices respectively.

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 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. 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 agents, 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 the match $\mu : N \to N$, where $N = W \cup F$, describes which worker is matched to which firm.[12] There is generically a unique match that maximizes the total surplus generated in the market.[13] We denote this efficient match by $\mu^* : N \to N$, so that

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

3 Theoretical Predictions

To guide our experimental investigation, we outline four prominent theories proposed in the literature. These theories yield different predictions of players' expected payoffs

[12]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.

[13]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.

and matches—and thus of 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 Appendix A for additional details. We chose Games 15, 25, and 30 to help us differentiate between these theories—they make very different predictions about final outcomes in these markets, and how these outcomes change across the games.

Theoretical predictions in our environment fall into two categories: those that use cooperative game theory and those that use non-cooperative game theory. The two cooperative theories we will consider are (i) an extension of Nash bargaining where the disagreement points are defined endogenously by the market; and (ii) the mid-point of the core. Both of these theories have been extensively applied in our setting. 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 agents is required to split between themselves the surplus they generate.[14]

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. Of course, this depends on the agreements others have reached, and so the solution boils down to finding a fixed point of a large system of equations.

Worker i's disagreement payoff is given by $\underline{u}_i = \max\left(0, \max_{j \in F} s_{ij} - v_j\right)$, and the firms' disagreement payoffs are defined analogously. Given these disagreement payoffs, an outcome is an SPB outcome if and only if the match is efficient and the payoffs solve the following system of equations:

$$u_i = \underline{u}_i + \frac{1}{2}\left(s_{ij} - \underline{u}_i - \underline{v}_j\right) \qquad \text{for all workers } i \tag{1}$$

$$v_j = \underline{v}_j + \frac{1}{2}\left(s_{ij} - \underline{u}_i - \underline{v}_j\right) \qquad \text{for all firms } j \tag{2}$$

It is shown in Rochford (1984) that a solution to this system of equations always exists. While in principle there can be multiple solutions in all the games we study it is unique.

Remark 1. *In all the games we study there is a unique symmetric pairwise bargained*

[14] Suppose, for example, there is one worker and two firms and that all matches generate a positive surplus. The Shapley value will then require that the match generating the higher surplus be implemented but that the unmatched firm nevertheless receive a strictly positive payoff. We view the Shapley value as a better solution concept for determining how surplus is split within a fixed coalition, rather than for studying which coalitions—or in our case, which matches—form.

outcome, and this outcome corresponds to (i) the nucleolus, (ii) the kernel, and (iii) the pre-kernel.[15]

While we relegate a derivation to Appendix A.1, solving this system of equations for Games 15, 25, and 30 yields the predictions shown in Table 2.

Table 2: Theoretical predictions of the symmetric pairwise bargained outcomes

	Game 15			Game 25			Game 30		
	eff.	B (C)	A (D)	eff.	B (C)	A (D)	eff.	B (C)	A (D)
SPB	100%	8.7	11.3	100%	5	15	100%	3.3	16.7

The symmetric pairwise bargained outcomes represent one selection from the core. Our next cooperative theory is an alternative refinement of the core. The core comprises the set of bargaining outcomes that are *feasible* such that the payoffs of all workers and firms sum to weakly less than the total surplus generated across all matches, and for which there is no coalition of workers and firms that could find an alternative 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. In other words, a bargaining outcome is in the core if and only if it is *pairwise stable*, which requires that for all workers i and all firms j, $u_i + v_j \geq s_{ij}$.

Shapley and Shubik (1972) also show that in all pairwise stable/core outcomes, the match that is implemented must maximize the total surplus. As generically there is a unique match with this property, pairwise stability alone pins down who must be matched to whom and there is no scope for inefficiency. While pairwise stability, or equivalently the core, pins down the match, many payoff vectors can typically be supported as core outcomes. In particular, Shapley and Shubik (1972) show that there is a core outcome in which all agents on one side of the market simultaneously receive their minimum possible core payoff, while all agents on the other side of the market simultaneously receive their maximum possible core payoff.

An alternative refinement of the core to the symmetric pairwise bargained outcomes uses the payoff structure of the core identified by Shapley and Shubik (1972). Either the extreme points can be used (see, for example, Kranton and Minehart (2001)) or some convex combination of the extreme points can be used (Corominas-Bosch (2004),Elliott (2015)).[16] As in our experiment there is no difference between the two sides of the market, we look for the **mid-point of the core** (mid-point).

[15]Equivalence of the kernel, pre-kernel, and SPB outcomes holds generally for all assignment games and follows from results in Rochford (1984) and Driessen (1998). The nucleolus is contained in the kernel, and so equivalence of it with the other solution concepts follows the uniqueness of the SPB outcomes for the games we study.

[16]Corominas-Bosch (2004) studies a non-cooperative game that selects the different convex combi-

In a non-transferable utility (NTU) environment, Enchenique and Yariv (2013) find experimentally that the median stable match is reached. For the markets we consider, the mid-point of the core corresponds to the transferable utility (TU) median stable matching (Schwarz and Yenmez, 2011).

Redefining $\underline{\mathbf{u}}'$ as the workers' minimum core payoffs, letting $\overline{\mathbf{u}}'$ be the workers' maximum core payoffs, redefining $\underline{\mathbf{v}}'$ as the firms' minimum core payoffs, and letting $\overline{\mathbf{v}}'$ be the firms' maximum core payoffs, the bargaining outcomes $(\underline{\mathbf{u}}', \overline{\mathbf{v}}', \mu^*)$ and the bargaining outcomes $(\overline{\mathbf{u}}', \underline{\mathbf{v}}', \mu^*)$ are in the core. Moreover, as the core is convex, the mid-point of these outcomes, $\left(\frac{1}{2}(\underline{\mathbf{u}}' + \overline{\mathbf{u}}'), \frac{1}{2}(\underline{\mathbf{v}}' + \overline{\mathbf{v}}'), \mu^*\right)$, is also in the core. As $u_i + v_{\mu^*(i)} = s_{i\mu^*(i)}$, these outcomes simplify to

$$u_i = \underline{u}_i' + \frac{1}{2}\left(s_{ij} - \underline{u}_i' - \underline{v}_j'\right) \qquad \text{for all workers } i \tag{3}$$

$$v_j = \underline{v}_j' + \frac{1}{2}\left(s_{ij} - \underline{u}_i' - \underline{v}_j'\right) \qquad \text{for all firms } j \tag{4}$$

This is the same payoff structure that we found in Equations (1) and (2), but with each disagreement payoff replaced by the minimum payoff that player could receive in any core outcome. In the four-player games we study, $\underline{u}_A' = \underline{v}_C' = 0$. In Game 15, $\underline{u}_B' = \underline{v}_D' = 0$; in Game 25, $\underline{u}_B' = \underline{v}_D' = 5$; and in Game 30, $\underline{u}_B' = \underline{v}_D' = 10$. This generates the predictions recorded in Table 3.

Table 3: Theoretical predictions of the core and mid-point of the core

	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]
Mid-Point	100%	10	10	100%	7.5	12.5	100%	5	15

For completeness, and despite making only set-valued predictions of the payoffs, we also include the range of **core payoffs** each player can receive.[17] All the cooperative theories we consider capture the idea that the simple threat of players A and D reaching agreement and leaving B and C unmatched should be enough to induce B and C to reach agreements that do not leave A and D with a profitable deviation.

The two remaining theories we consider use non-cooperative game theory and thus require the specification of a bargaining protocol. We use a standard protocol which extends Rubinstein bargaining to accommodate many players, and which we view

nations of the extreme points of the core as an exogenous parameter is varied, but in an environment in which the gains from trade are either 1 or 0.

[17] Some papers, such as Cole et al. (2001a), look for results that are robust to any selection from the core.

as natural for the environments we consider. It is also closely related to multilateral bargaining protocols (à là Baron and Ferejohn (1989)). A key feature is that we allow players to choose whom to make offers to instead of modeling random meetings. This prevents players from having to sometimes delay in order to match efficiently, and for this reason we view it as giving efficient outcomes a better shot than the alternative.

The 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 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.

There are potentially multiple equilibria of this dynamic game. We will focus on two types of equilibria that are prominent in the literature. A first type of equilibria we will look at are those that are subgame perfect and involve stationary strategies, insofar as strategies can depend only on who is left in the market rather than the entire history of play. These are the **Markov perfect equilibria** (MPE) where the state of the world is given by the set of players who are active and have not yet exited the market. We will see that the Markov perfect equilibria predict that matches will sometimes be inefficient. However, efficient outcomes can be obtained by non-stationary strategies that are history dependent. A second type of equilibria we will consider are efficient perfect equilibria (PE). We will construct these equilibria for the games we analyze. Subtle rewards and punishments can be required, and to the best of our knowledge known folk theorems do not apply in this setting. Crucially, once players exit, they cannot be punished or rewarded.

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). Many papers in the bargaining literature, 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), focus on the MPE. The exact bargaining protocol these papers study differ. Our setup corresponds to that in Elliott and Nava (2015)—players choose whom to make offers to and exit the market upon reaching agreement.

Only in Game 15 does there exist an efficient MPE, since the link between players A and D does not bind and can affect the players' bargaining positions. On the contrary, Games 25 and 30 often end in an inefficient outcome, with the likelihood of inefficiency increasing in the value of the diagonal link. Some intuition for this result comes from the observation that the diagonal link serves as an outside option for players A and D. For this link to "matter" (affect bargaining outcomes of players), it has to be exercised with positive probability on path, which results in an inefficient outcome for the whole market. This differs from the case of bilateral bargaining with an exogenous outside option. In that case a binding outside option can be exercised with vanishingly small probability and still discipline payoffs. The difference here is that the outside option does not last forever. For example, if players A and C exit the market, then player D will end up bargaining bilaterally with B without an alternative. It is this temporary nature of alternative possible matches that results in players having to exercise them with strictly positive probability to affect their payoffs.

Here we provide a brief and rough guide to deriving the limit MPE payoffs; for a full derivation the reader is referred to Appendix A.2. Solving the game backwards: If the first match reached is a core match, then the unmatched players will bargain bilaterally and receive limit payoffs of 10 each. If A and D reach agreement first, then B and C will be left unmatched and receive a payoff of 0. 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. In Game 25, efficient proposals are no longer an equilibrium. Suppose A always offered to C and D always offered to B. Then C could wait for D and B to reach agreement first, and then bargain bilaterally with A, receiving a limit payoff of 10. Similarly, B must get a limit payoff of at least 10. But if B and C each get at least 10, then A and D each get at most 10 apiece and would have a profitable deviation to offer to each other. In equilibrium A and D mix between offering to each other and offering to their efficient partners. Doing so depresses the expected payoffs of C and B, because they are left unmatched whenever A and D reach agreement. Indeed, in equilibrium A and D mix to the extent that they reduce the continuation values of B and C just enough for it to be equally profitable for them to make an efficient offer as to offer to each other. Whenever A and D reach agreement the total surplus generated in the market is 25 instead of 40 and there is inefficiency.

In Game 30, A and D reach the corner solution in which they offer to each other with probability 1. They still accept offers from their efficient partners, but whenever they are selected as the proposer there is mismatch. Thus as the value of the diagonal link AD increases, the frequency with which players A and D propose efficiently decreases, and so does overall efficiency. In the MPE there is a monotonic decrease in overall ex-ante efficiency as the value of the diagonal link increases from Game 15 to Game 25 to Game 30. The MPE predictions are summarized in Table 4.

Allowing players to use non-Markovian strategies enlarges the set of equilibrium outcomes. In particular, by using history-dependent strategies one can reach efficient

Table 4: Theoretical predictions of Markov perfect equilibrium

	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

Notes: In this table, we list limiting payoffs of unique MPE in each game as players become infinitely patient ($\delta \to 1$).

outcomes in all our games. This reflects a result by Abreu and Manea (2012a). 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. While in many games reaching efficient outcomes may involve complicated strategies, in Game 25 this is not the case. A full derivation of these equilibria is relegated to Appendix A.3; here we provide only an overview.

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. Thus once this subgame is entered there is no scope for rewards or punishments.

In Game 25 there is an efficient PE in which the strong players (A and D) can use the threat of offering to each other to ensure that an efficient outcome is played. We label these outcomes **efficient perfect equilibria with Markov reversion** (Eff. PE (i)). The threat of reverting to playing the MPE in which they mix between offering to each other and offering efficiently is credible, because the MPE are perfect equilibria. This threat is sufficient to ensure that the weak players accept offers weakly greater than their MPE payoffs. Nevertheless, after such an offer (from a strong player to a weak player) is accepted the remaining unmatched pair bargain bilaterally and the weak player receives a limit payoff of 10. Thus in this equilibrium, in the limit weak players receive their expected MPE payoffs with probability 1/2 and a payoff of 10 with probability 1/2. This means that the weak players appropriate all the efficiency gains and more while the strong players end up with an expected payoff less than their expected MPE payoff. Nevertheless, strong players do not have a profitable deviation when proposing. Conditional on being the proposer, a strong player's expected payoff is well above their expected MPE payoff.

Constructing an efficient PE in Game 30 is more complicated. MPE reversion does not provide strong enough incentives for the strong players to optimally offer to their efficient partners, but more complicated strategies can be used. In Appendix A.3 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. While we do not rule out the possibility that more complicated strategies can support a wider range of payoffs in a perfect equilibrium than the strategies we construct, we can place a lower bound on the expected payoffs that the weak players B and C must receive in expectation in any efficient perfect equilibrium.

Proposition 1. *In Game 25 and Game 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.*

Proof. In any efficient perfect equilibrium either B or C will be left to bargain bilaterally with their efficient partner. By efficiency, the limit surplus generated in any such subgame is 20. Moreover, there is a unique perfect equilibrium of any such subgame in which the remaining players receive equal shares. Thus in all efficient perfect equilibria either player B or player C must receive a limit payoff of 10 and so in expectation the sum of player B and player C's limit payoffs must weakly exceed 10. □

Table 5 summarizes the theoretical predictions of the theories discussed above in terms of final outcomes: the frequency with which an efficient match is reached, and players' payoffs by their network position. For the efficient perfect equilibrium without MPE reversion, we report the range of payoffs that can be supported by the strategies we construct.

Table 5: 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)
Coop.									
SPB	100%	8.7	11.3	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
Non-Coop.									
MPE	100%	10	10	72%	6.45	11.45	50%	4.17	13.33
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%	$(6\frac{1}{9}, 9\frac{4}{9})$	$(10\frac{5}{9}, 13\frac{8}{9})$

Notes: In the MPE row, we list limiting expected payoffs of players as $\delta \to 1$. For efficient PE, we consider two specifications: In (i) there is MPE reversion following a deviation, while in (ii) there are two off-path punishment states: one to punish A and B while rewarding C and D, and another to punish C and D while rewarding A and B.

4 Experimental Design and Procedures

We conducted our experiments 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.[18,19] 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.

As discussed in Section 3, we ran three treatments (Exit 15, Exit 25, and Exit 30) corresponding to the three markets described in Figure 1. 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 learning is natural to expect.

Within each game, we implement a bargaining protocol strategically equivalent to the one described in the previous section. 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 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 random termination of a 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. 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

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

[19]In the Supplementary Appendix, Section 2.1, 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.

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, Section 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 the 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.[20] Two features of our interfaces 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. If anything, this design choice should make it easier for the players to play non-Markovian strategies and reach efficient outcomes.

5 Results

In this section we present the results of the experiments described in the previous section. First, we run the horse race between the theories discussed in Section 3 and evaluate their performance based on the observed outcomes. Second, we dig deeper into predictions of the theory which emerges from this horse race as a clear winner. We explore market dynamics and the evolution of players' bargaining power before they reach the final outcomes, as well as individual strategies used by our experimental subjects.

5.1 Approach to Data Analysis

Unless otherwise noted, our results compare the average outcomes of two groups. 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 a random-effects GLS regression in which we regress the variable of interest (i.e., the payoffs of the players or an indicator of

[20]The list of questions and the screenshots of the game are presented in the Supplementary Appendix, Sections 1.2 and 1.3.

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.[21] We conclude that there is a significant difference between the outcomes of the two groups under consideration if the estimated coefficient on the group indicator dummy variable 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.[22] 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.

Finally, we note that while our 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.

5.2 Market Outcomes

Table 6 summarizes observed and predicted characteristics of the final match in each market. To remove learning effects we focus on the second half of each experimental session (last five repetitions). Figure 2 depicts the evolution of final match efficiency for each market separately.

We observe a significant decrease in efficiency levels as the value of the diagonal link increases from Game 15 to Game 25 to Game 30. Indeed, while all the final matches in the last five repetitions of Game 15 are efficient, efficiency drops to 51% in Game 25 and even further (to 30%) in Game 30. Regression analysis confirms that this monotonic decrease in efficiency is significant at the 5% level ($p < 0.01$ for Game 15 vs. Game 25, and $p = 0.01$ for Game 25 vs Game 30).[23] Figure 2 shows that from

[21] Recall that in our experiments subjects were re-matched to form new groups between games, thus, creating interdependencies between all observations in the same session. Running regression analysis with clustering standard errors by session accounts for such interdependencies.

[22] Recall that natural completion of the game occurs when any two players from the subset of players that propose new matches in the current round could generate a surplus of 0, while random termination occurs with probability 1% in each round irrespective of the participants' behavior.

[23] As described in the data analysis section above (Section 5.1), we run two separate regressions. The first regression compares efficiency in Games 15 and 25. Specifically, we run random-effects GLS regression in which we regress indicator of efficient final match in the last 5 repetitions of the sessions on the dummy variable that indicates Game 25 and a constant, clustering standard errors by session. The estimated coefficient on the dummy variable is negative and we report p-value associated with

Table 6: Predicted versus observed outcomes, last five repetitions

	Game 15			Game 25			Game 30		
	eff.	B (C)	A (D)	eff.	B (C)	A (D)	eff.	B (C)	A (D)
Coop.									
SPB	100%	8.7	11.3	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]
Non-Coop.									
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: In the last two rows, we 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.

the 4th repetition onwards the efficiency of the final matches decreases monotonically from Game 15 to Game 25 to Game 30. The columns headed eff. in Table 6 show the predicted frequency of the efficient matches for the different theories and the actual frequency of the efficient matches. The only theory that predicts a monotonic decrease in efficiency levels over our three treatments is Markov perfect equilibrium. All the other theories predict full efficiency irrespective of the underlying game—a prediction which is clearly rejected by our data.

We now turn our attention to players' payoffs, distinguishing between those with one link (players B and C), whom we refer to as the weak players, and those with two links (players A and D), whom we refer to as the strong players. The next-to-last row in Table 6 reports the average payoffs of the players observed in the second half of the experiment, along with the corresponding standard errors. The final row reports the average payoffs, conditional on an efficient match being reached.

All the theories either predict or are at least consistent with a monotonic increase (decrease) in the payoffs of strong (weak) players as we move from Game 15 to Game 25 to Game 30.[24] Our data confirm this prediction. The value of the diagonal links directly translates into the bargaining power of strong players, as they obtain significantly higher payoffs in Game 30 than in Game 25, and in Game 25 than

that estimated coefficient. The second regression performs similar analysis for comparing efficiency of final matches in the last 5 repetitions of Games 25 and 30.

[24]Even the theories that make set-valued predictions (the core and efficient PE (ii)) exhibit monotonic changes in the range of payoffs that can be supported.

Figure 2: Evolution of final match efficiency, by market

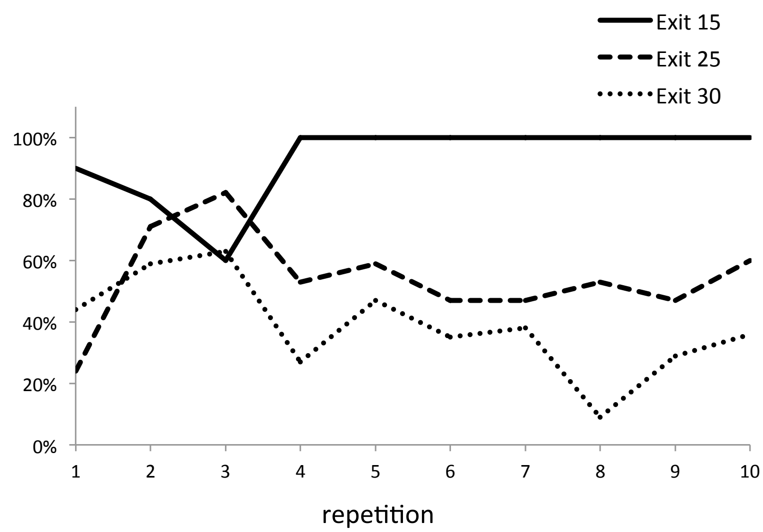

in Game 15. At the same time weak players obtain significantly lower payoffs in Game 30 than in Game 25, and in Game 25 than in Game 15, as their bargaining positions are weakened by the high-value diagonal link connecting the two strong players ($p < 0.01$ in all pairwise comparisons). All the theories also predict that the payoffs of the strong players are weakly higher than those of the weak players. This prediction is borne out in our data. In Game 15 players' payoffs do not depend on their network position ($p = 0.319$), which is predicted by all the theories except symmetric pairwise bargaining. In both Game 25 and Game 30 we find that strong players receive significantly higher payoffs than weak players ($p < 0.01$), as predicted by all the theories. This observation confirms that our experimental subjects have a basic understanding of the bargaining games they are facing.

Despite the agreement of our data with the above comparative static predictions, none of the theories make point predictions that match payoffs very closely. In Game 15 all the theories except symmetric pairwise bargaining do well, but 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, and possibly a larger range than those that the strategies we construct can support, there does not exist an efficient perfect equilibrium that closely matches the observed payoffs in Game 30. Proposition 1 showed that 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 (see Table 6).

The final payoffs of players depend on whether or not an efficient match was reached, since mismatches automatically mean payoffs of 0 for the weak players. Given the massive inefficiencies observed in Games 25 and 30 and the fact that the only theory that predicts mismatches is MPE, this puts the other theories at a disadvantages in terms of their ability to predict players' payoffs. 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 6). For comparison, we also present MPE-predicted payoffs of players who reached efficient outcomes (fifth row of Table 6).

A few interesting patterns emerge with respect to the payoffs of players when they reach an efficient outcome. First, the payoffs of strong players still monotonically increase from Game 15 to Game 25 to Game 30, while the payoffs of weak players still monotonically decrease from Game 30 to Game 25 to Game 15. This is consistent with the predictions of all the theories, including MPE conditional on efficient matches being reached. The point predictions made by symmetric pairwise bargaining and mid-point of the core are still rejected by the data, while the observed payoffs fall inside the wide range of payoffs that can be supported by the core. With respect to non-cooperative theories, MPE predicts players' payoffs within a 95% confidence interval in Game 25, and it comes close to doing so in Game 30 but still falls outside the 95% confidence interval. Finally, there exist efficient PE strategies that are able to generate the observed payoffs in both games.

We finish this section by concluding that, overall, MPE emerge as a clear winner among the theories considered in terms of organizing market outcomes in the two dimensions: efficiency of final matches and players' payoffs. On the payoff dimension alone it is hard to distinguish one theory from another, but MPE does as well as any other theory, and on the efficiency dimension it does much better than any other theory. Importantly, we find that our markets often fail to clear efficiently and mismatches are common.

5.3 Other Predictions of Markov PE

To better understand what drives the extensive mismatch we find, we now consider some more refined MPE predictions regarding market dynamics and the evolution of bargaining power. The main way in which the MPE outperforms the other theoretical predictions is by more closely matching the frequency of mismatch and mismatch could, in principle, be driven by mistakes orthogonal to all the theories. The analysis in this section also addresses this concern.

5.3.1 Efficiency and the Network Position of the First Mover

In all the theories except MPE an efficient agreement is reached regardless of the identity of the first mover. According to MPE, whether or not the market clears efficiently in Games 25 and 30 depends on which player was randomly selected to make the first move. If the first mover is a strong player (a player with two links), then the market is predicted to end up with an inefficient outcome with positive probability, while if a weak player (a player with one link) moves first it is predicted that an efficient outcome will be reached with certainty. This pattern follows directly from two observations: (a) On the equilibrium path all the players make offers that are accepted, and so there is immediate agreement, and (b) the strong players propose to each other with positive probability in Games 25 and 30. In Game 15, by contrast, efficiency is predicted to always be reached irrespective of the identity of the first mover, which is consistent with our data, as we observe that all the final matches in the last five repetitions of Game 15 are efficient.

Table 7: Effect of network position of the first mover on final match efficiency, last five repetitions

	Game 25		Game 30	
	Regression (1)	Regression (2)	Regression (1)	Regression(2)
First mover is strong	-0.37^{**} (0.11)		-0.49^{**} (0.10)	
First accepted offer by strong player		-0.81^{**} (0.07)		-0.93^{**} (0.04)
Constant	0.66^{**} (0.07)	0.97^{**}(0.06)	0.53^{**} (0.06)	0.95^{**} (0.04)
# of observations	81	81	69	69
# of sessions	4	4	3	3
overall R-sq	0.1301	0.6242	0.2853	0.8678

Notes: Random-effects GLS regressions. Dependent variable is an indicator of an efficient final match. Standard errors are clustered at the session level. ** indicates significance at 5% level.

Table 7 reports the results of a regression analysis in which we study the link between the network position of the first mover and the final match efficiency. For each game, we run two regressions; in both of them, the dependent variable is an indicator for the efficiency of the final matches. In regression (1) the right-hand side variables include a binary variable that takes value 1 if the first mover is a strong player and 0 otherwise. In regression (2) the right-hand side variables include a binary variable that takes value 1 if the first accepted offer was made by a strong player and 0 otherwise. While MPE predicts that on path no offers should ever be rejected, such behavior is present in our data, which is why we also run regression (2).[25]

Consistent with the MPE predictions, we observe that in both Game 25 and Game

[25] We refer the reader to Section 5.3.3 for a discussion of the rejected offers.

30 efficiency is significantly lower in games in which the first mover is a strong player rather than a weak player. The same is true if we condition the efficiency of matches on the network position of the player who makes the first accepted offer. In regression (2) the higher estimated coefficient in Game 25 than in Game 30 is consistent with the MPE predictions regarding frequencies of efficient proposals by strong players in these two games: In Game 25 strong players are supposed to mix between proposing efficiently and not, while in Game 30 they should always propose inefficiently. We come back to this point in the next section, when we discuss the strategies used by our experimental subjects.

5.3.2 Payoffs Conditional on an Efficient Match by Market Composition

Markov Perfect Equilibrium predicts that the payoffs of players conditional on reaching an efficient match will depend on the composition of the market when they reach agreement.[26] When all players are active the strong players (those with two links) have a stronger bargaining position than the weak players (those with one link), because of the additional trading opportunities that they have with each other. In contrast, once one pair comprised of a strong player and a weak player reach agreement the remaining strong player and weak player are left to bargain bilaterally and have equally strong bargaining positions. Thus, conditional on reaching an efficient match, the payoff of the second weak player to match is predicted by MPE to be equal to the payoff of the second strong player to reach agreement, while the payoff of the first weak player to match is predicted to be substantially less than the payoff of the first strong player to match. Equally, conditional on reaching an efficient match, weak players receive a higher payoff when they match second, while strong players receive a higher payoff when they match first. We test these predictions.

Figure 3 presents histograms of the final payoffs of the strong players by their order of exit and conditional on an efficient outcome being reached. In Figure 3 grey bars depict payoffs of players that exited the market first, while black bars represent payoffs of players that exited the market second. Only the payoffs of strong players are shown, but as we are conditioning on an efficient match being reached, the payoffs of weak players are 20 less the payoff of the strong players. As is evident from Figure 3, strong players enjoy a premium for exiting the market first, while, correspondingly, weak players receive lower payoffs when they exit first. In Game 25, the average payoff of strong players is 12.4 if they exited first, while it is only 10.1 if they exited second. For weak players it is 7.6 if they exited first and 9.9 if they exited second. Similarly, in Game 30 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

[26] We condition on an efficient match to separate out the effect of market composition on payoffs from efficiency considerations. Similar results hold when we do not condition on an efficient match being reached.

Figure 3: Average payoff of the next strong player to reach an agreement with her efficient partner by market composition (whether all players are active and unmatched or one pair has already exited)

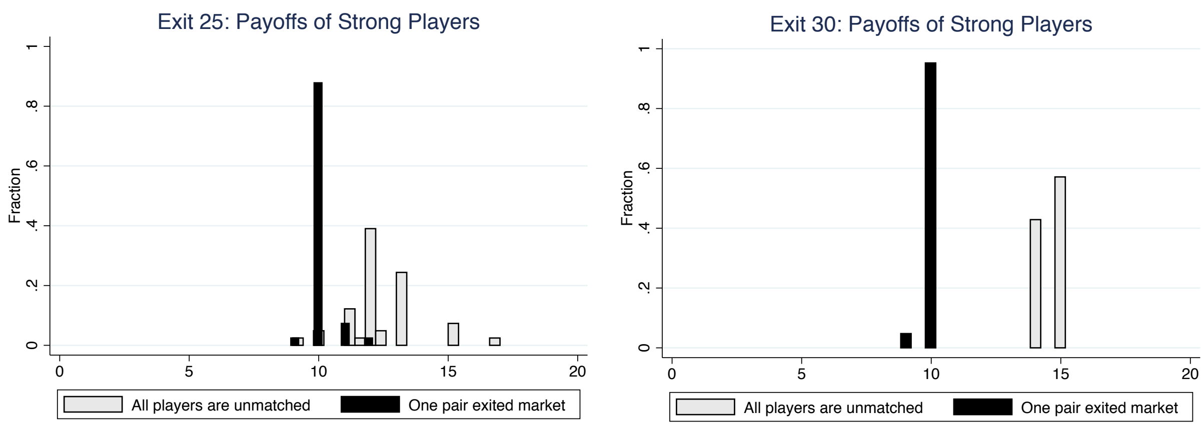

Notes: In both figures, we consider only groups that reached efficient match and only the last five repetitions.

compared to 10 when exiting second. These differences are statistically significant in a regression analysis ($p < 0.01$ in both treatments).[27]

A natural question is whether our experimental subjects tried to manipulate the order of exit in their favor. Given the parameters of our games, the unique MPE predicts that all subjects should always propose and never delay. Weak players are incentivized to make offers by the possibility of the strong players reaching agreement with each other and exiting the market, leaving them unmatched. In other words, it is because an inefficient match is sometimes implemented that they should not want to delay, hoping that they will get to bargain bilaterally with their efficient partner. One way to investigate this issue is to consider how often subjects chose not to make a proposal, which they could do by pressing the "Do Nothing" button. In general, we observe very few delays (less than 2.5% in both treatments). Moreover, regression analysis detects no significant difference between the frequency of delays between strong and weak players conditional on the market being complete ($p > 0.10$ in both Game 25 and Game 30). This finding is consistent with the MPE predictions and suggests that weak subjects (those with one link) understand the costs of delay relative to its benefits.

[27]For comparison, the average payoff of a strong player that matches inefficiently is 12.5 in Game 25 and 15 in Game 30.

5.3.3 Players' Strategies

To further investigate MPE predictions we ask whether the play of players is consistent with 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 allows players' strategies to depend on the history of play, but only as reflected by the set of players who remain in the market (the state). This greatly simplifies the strategy space, and, given the structures of the markets we consider in this paper, 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.

In the second state, when only two players remain in the market, there is a unique perfect equilibrium, in which each of the two unmatched players offers half of the available surplus (10 out of 20) and accepts any offer that guarantees at least that amount. Looking at all the proposals submitted by the players when the markets were incomplete, we observe that, consistent with this theoretical prediction, many of them offered exactly half of the surplus. In particular, in the last five repetitions of the Game 15 treatment, 47% of all proposals submitted for incomplete markets were requests to keep 10 out of 20, while in the other two treatments these fractions are even higher, reaching 58% and 83% in the Game 25 and the Game 30 treatments, respectively. The remaining proposals were generally above those predicted.[28] In terms of acceptance behavior, in these states the vast majority of proposals that guaranteed a player a payoff of at least 10 were accepted in all three treatments.[29]

In the remainder of this section we focus on players' strategies when all the players are active. We begin by considering the relative frequency with which strong players offer inefficiently to other strong players. Markov Perfect Equilibrium predict that strong players will never offer to each other in Game 15, sometimes offer to the other strong player in Game 25, and always offer to the other strong player in Game 30. In the Game 25 treatment, given that only a few offers are observed for each strong player, mixing as prescribed by MPE would result in a wide range of players' empirical mixing ratios.

In Figure 4 we present the empirical histogram of individual frequencies of proposing efficiently in the last five repetitions in each session. In this figure, we also plot (in

[28]Average offers of the unmatched strong and weak players left them with 11.7 and 11.3 in the Game 15 treatment, 11.5 and 10.3 in the Game 25 treatment, and 10.5 and 10.3 in the Game 30 treatment, respectively.

[29]In particular, in the Game 15 treatment this fraction is 76%, while it is 95% in the Game 25 treatment and 100% in the Game 30 treatment. These results are for the last five repetitions in each session.

Figure 4: Frequency of efficient proposals by strong players, last five repetitions

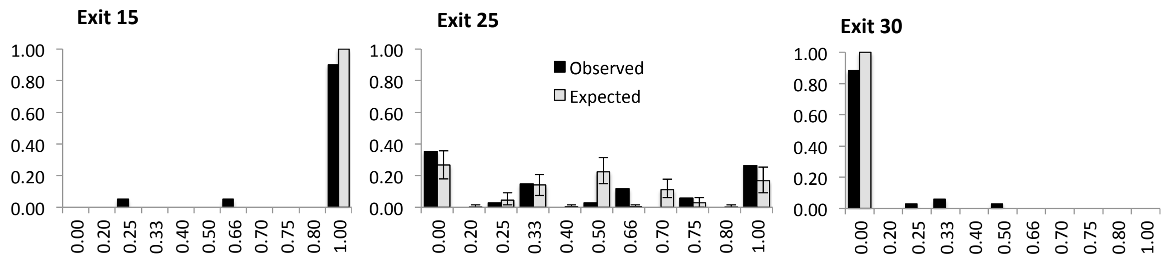

Notes: For each subject, we compute the frequency of proposing to her efficient partner in the first round of the last five repetitions conditional on this player being assigned a strong position. Black bars are the observed frequencies, while light grey bars represent the expected frequencies if all players were to play exactly the strategy prescribed by MPE. For Game 25, we calculated expected frequencies using the number of observations we have in this treatment, along with the 95% confidence intervals.

grey bars) the expected ratios for each game were all players to play the MPE strategy, along with the 95% confidence intervals based on simulations given the number of data points we have in our experimental sessions. As in the MPE strategy agreements are reached immediately, we restrict attention to the first round to analyze how this aspect of the theory fits the data. We consider the dynamic play later, but note that including all rounds in which markets were complete generates a figure very similar to Figure 4 (see the discussion in the Supplementary Appendix, Section 2.2). Consistent with MPE prediction, we observe a significant shift in strategies used by the strong players as the value of the diagonal link increases. In particular, in Game 15 over 80% of subjects always propose to their efficient partner in the last five repetitions, while in Game 30 over 80% of subjects never propose to their efficient partner in the last five repetitions. The observed mixing in Game 25 also fits the theory well. Statistical analysis confirms that subjects are significantly more likely to make efficient proposals in Game 15 than in Game 25 ($p < 0.01$), and in Game 25 than in Game 30 ($p < 0.01$), as predicted by MPE.

We turn now to the amounts players wanted to keep when making offers. Table 8 describes players' predicted play according to the MPE strategies. For completeness, we include the mixing probability as well as the amount offered. Since on the equilibrium path players should accept equilibrium proposals made by other players, this table automatically summarizes equilibrium acceptance behavior for each network position.

Figure 5 depicts box plots of average absolute differences between amounts offered in Game 25 and Game 30 and those predicted by MPE, by offer type (strong player to strong player, weak player to strong player, etc.) and constructed for the first five and the last five repetitions in each session separately. The differences were first computed

Table 8: Markov perfect equilibrium strategies, by game

	Game 15		Game 25		Game 30	
	freq	ask amount	freq	ask amount	freq	ask amount
Strong to Strong	0%	–	56%	13.55	100%	15.83
Strong to Weak	100%	10	44%	13.55	0%	–
Weak to Strong	100%	10	100%	8.55	100%	6.67

and averaged for each subject, and then combined with those of other subjects. Each of the figures in Figure 5 corresponds to one of the offers on the equilibrium path, as indicated in Table 8. In these figures, each box shows the interquartile range (between the 25th and 75th percentiles), with the median value indicated by the thick dashed line. The length of whiskers is set at the standard 1.5 times the interquartile range.

These offers deviate from those predicted by MPE, but not too substantially. Moreover, as subjects gain more experience with the game, they play strategies which more closely resemble the ones predicted by MPE. In all cases there is a significant decrease in the average absolute difference between amounts offered and those predicted by MPE. To reach this conclusion we have regressed absolute differences between amounts offered and Markov predictions on the dummy variable that indicates the last five repetitions in each session, while clustering observations by the session.[30] We find that the estimated dummy coefficient is negative and different from 0 at the standard 5% significance level in all offers depicted in Figure 5.

Finally, we consider which offers subjects accept and which they reject, and how that varies with their network position. Figure 6 depicts accepted and rejected offers in each treatment in the last five repetitions for each network position as well as the number of observations we have for each offer type. As is evident from this figure, responders' behavior closely mimics the proposing behavior of subjects: As the value of the diagonal link increases, strong responders accept only higher offers, while weak responders are willing to accept smaller shares. Regression analysis confirms this pattern: Shares accepted by strong responders are significantly higher in Game 25 than in Game 15 ($p < 0.01$) and also significantly higher in Game 30 than in Game 25 ($p < 0.01$). Similarly, shares accepted by weak responders are significantly lower in Game 25 than in Game 15 ($p < 0.01$) and also significantly lower in Game 30 than in Game 25 ($p < 0.01$), although we have only a small number of observations in this last case since strong responders almost exclusively propose to each other in this treatment (see Figure 4).

The evidence we have discussed, and the interpretation we have given to it, is somewhat selective. We have taken evidence that players get closer to playing in line with the MPE strategy over the duration of the experiment as a positive for

[30]The results do not change if we cluster standard errors at the individual level.

Figure 5: Average absolute deviations of the amounts offered by players from the MPE predictions

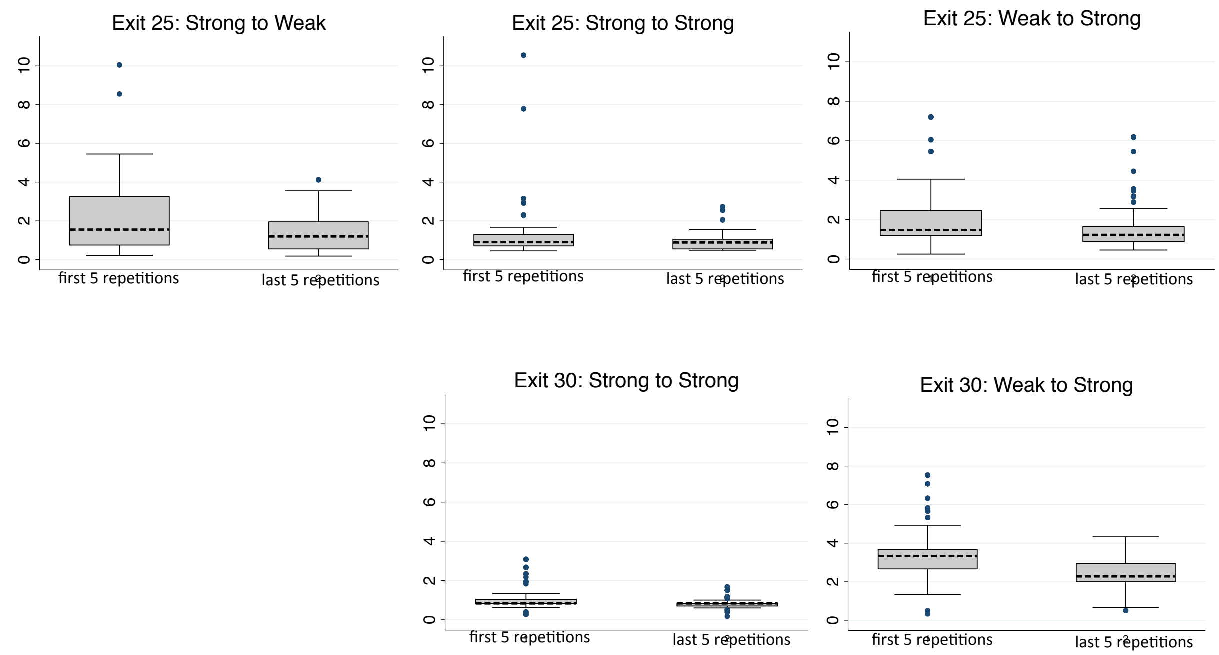

Notes: Averages are computed separately for each subject in the first and last five repetitions of a session, and then combined with those of the other subjects. We focus only on cases in which markets were complete. Each box depicts the interquartile range (between the 25th and 75th percentiles), with the median value indicated by the thick dashed line. The length of whiskers is set at 1.5 times the interquartile range.

the theory, while it is of course in direct violation of the Markovian assumption. We have also not reported direct violations of non-stationary strategies. Indeed, no subject plays in exactly the same way or even makes exactly the same offers whenever all the players are active. Although learning and/or random errors are enough to reject the hypothesis that play is exactly Markovian, we are instead interested in a much weaker test of players' strategies—in the last five repetitions, after learning has hopefully finished, do players systematically deviate from the MPE strategy in terms of the average offers they make and their acceptance strategies—and for the strong players, the frequency with which they make inefficient offers? While we find some small systematic deviations from the strategy, overall the main comparative static predictions made by MPE across our three treatments are preserved. Such an analysis is important, as it sheds light on how subjects reached different final outcomes in the three games we consider.

We conclude this section by tying together what we learned about individual strategies and efficiency patterns documented in Section 5.2. Recall that in both

Figure 6: Responders' behavior by network position, last five repetitions

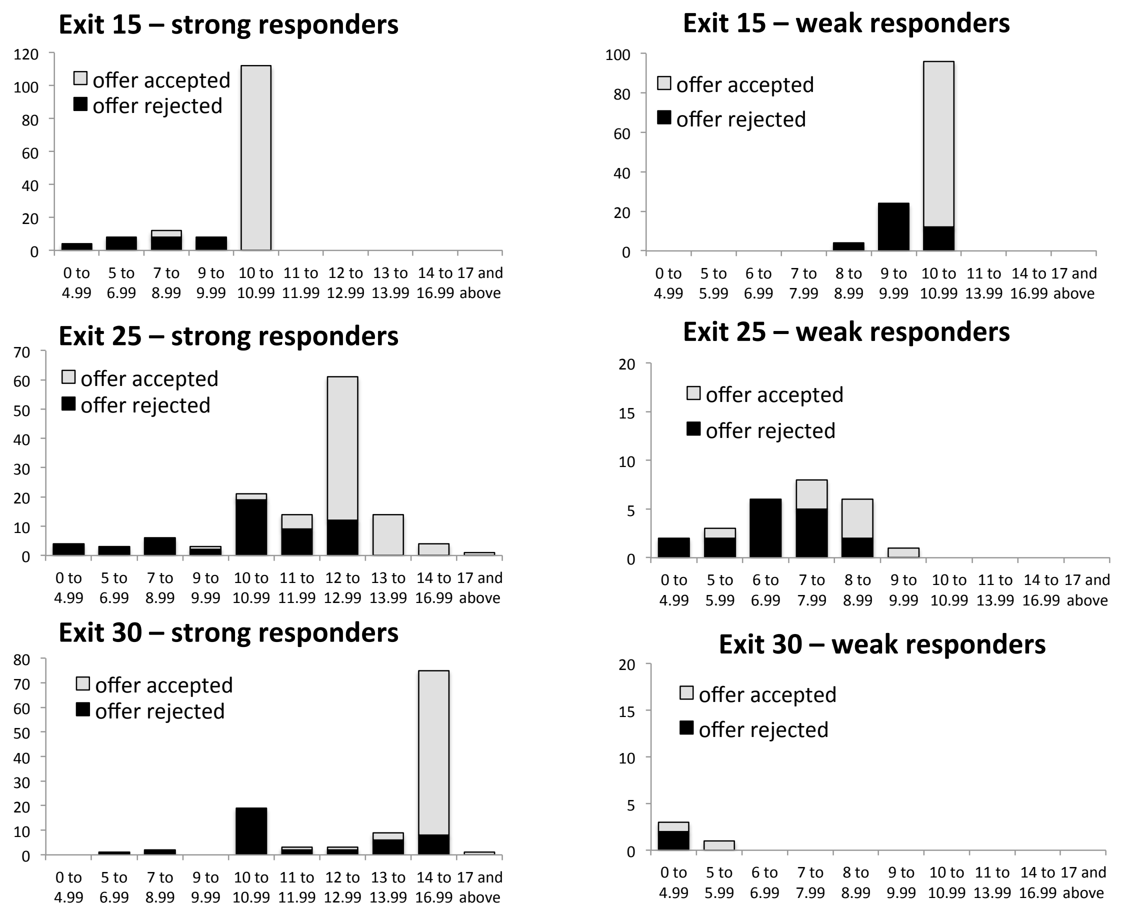

Notes: Offers received by responders are depicted on the horizontal axes. The height of each bar represents the number of observations in the indicated offer range. [[Marina, can we change the labels in the figure from offer to offers?]]

Game 25 and Game 30 we observe final efficiency levels which are significantly below those predicted by MPE, while in Game 15 this is not the case. To understand this pattern, we first note that there is no significant difference in the network position of the randomly selected first mover between any two treatments (using the Test of Proportions and the standard 5% significance level) [[Marina, is this phrasing right?]]. However, in Game 25 and Game 30 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, while there is no difference in Game 15. In particular, in Game 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 Game 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. This last feature, combined with observation that in both Game 25 and Game 30 strong players often make proposals to each other, as documented in Figure 4, explains why the observed efficiency levels are below those predicted by MPE in these two games.

6 Experiment II

Our most important findings are that the players match inefficiently in the Exit 25 and Exit 30 treatments, the frequencies of mismatch are increasing from the Exit 15 treatment to the Exit 25 treatment to the Exit 30 treatment, and the losses due to mismatch are substantial in the Exit 25 and Exit 30 treatments. In both cases more than 15% of the potential surplus is lost on average in the last five repetitions. These results are all consistent with the MPE, and analyzing the MPE might help us better understand what is driving these inefficiencies. Indeed, it is interesting to contrast the inefficiencies we find with those that would occur if two players bargained bilaterally with a long-lived outside option. In this case, in line with the "outside option principle" the two players would reach agreement with probability 1 and the player with the binding outside option would receive a payoff equal to its value. The crucial difference is that the outside options provided by alternative matches in the market are not long-lived. It is precisely because the market evolves and bargaining positions change that inefficiencies occur. If instead players were allowed to rematch, then alternatives would be long-lived and we might hope that efficiency is restored. Our second experiment investigates this possibility.

6.1 Design and Procedures

The design and experimental procedures of Experiment II are very similar to those of the main experiment reported in this paper. In this section, we briefly summarize the main features of Experiment II and refer the reader to the Supplementary Appendix, Section 1, in which we present instructions and screenshots for this experiment. We conducted Experiment II in the same two locations as the main experiment: at the ESSL at University of California, Irvine and at the EBEL at University of California, Santa Barbara. Ten sessions were conducted, 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.

Similar to the main experiment, we conducted three treatments (Stay 15, Stay 25, and Stay 30), with each treatment corresponding to one of the markets described in Figure 1. In each session, subjects played ten repetitions of the same game, with random rematching between games. The main feature of the Stay treatments is the possibility of reneging on agreements formed in previous rounds. Recall that in the Exit treatments players have no opportunity of reneging, as those who reach agreements are forced to exit the market permanently, which means they cannot make any further moves.[31] On the contrary, in the Stay treatments players who reach agreements do not exit the market and can unilaterally break agreements they are part of at a small cost c.[32] In all three Stay treatments, 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 matches proposed to her if such proposals come along. 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 as an unmatched player. 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 the Stay treatments were identical to those of the Exit treatments. In particular, just as in the Exit treatments, there were two ways in which a game could end in the Stay treatments. First, there was a 1% chance that the game ended after each round, determined by a random draw of the computer. Second, the game ended if there was no positive surplus remaining between any pair of players who have made proposals in the last round.

7 Experiment II: Theoretical Predictions

In this section we state a formal conjecture and provide some loose intuition. The intuition is not intended to take the place of a proof. We were unable to establish the result theoretically.

Conjecture 1. *For any market there exists $\overline{c} > 0$ such that for any cost of reneging $c < \overline{c}$ there exists an MPE that is efficient in the limit (i.e. achieves the maximum possible surplus as $\delta \to 1$).*

[31] 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.

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

Consider, for example, Game 30. Suppose players A and C have reached an agreement in which player A will receive a payoff of 15, and the remaining players players B and D are still unmatched. This situation is very similar to one in which B and D are bargaining bilaterally and D has an exogenous outside option with value $30-15-c = 15-c$. In this alternative game, there is a unique perfect equilibrium that yields the efficient match. If selected as the proposer, B would make an acceptable offer to D that would give D a limit payoff of $15-c$, while if selected as the proposer, D would makes an acceptable offer with probability 1 to B that also leaves D with a limit payoff of $15-c$. Solving now the game when no players are yet matched, it is consistent for the strong players to demand limit payoffs of 15 and for the weak players to demand limit payoffs of 5.

While the above reasoning is deceptively simple, establishing the conjecture formally is extremely challenging. To get a feeling for some of the issues that arise, first note that we need $\delta < 1$ to avoid the standard bargaining indeterminacy. But in this case, in the subgame in which A and C are matched there will be a strictly positive probability that A and D will reach agreement next (albeit a small probability that goes to 0 in the limit). Thus we need to specify what the players will do when A and D are matched. Furthermore, players' current agreements are payoff relevant for the other players, expanding the set of Markov states that need to be considered. As the sequence in which different players propose will affect the agreements reached, and as A and D always offer to each other with positive probability for $\delta < 1$, there can be infinite sequences of mismatches that occur before the game ends. This can blow up the set of Markov states that need to be considered. Finally, there is no guarantee that a player's continuation value is monotonic in the size of an offer they receive from another player. This inhibits the use of subgame perfection to pin down the value of offers than must be made.

While the intuition for why efficiency should be restored by allowing players to rematch is straightforward, establishing the result is hard. Moreover, it might be argued that if it is so hard to establish theoretically, we should not expect players to play in this way in practice. We therefore head straight to the lab to investigate what happens. Do the messages from the theoretical analysis of MPE with exit provide us with useful intuitions for thinking about what will happen without exit as well? Are conjectures based loosely on this kind of reasoning enough to guide our thinking in a useful way?

8 Experiment II: Results

In this section we report several key comparisons of final outcomes and the strategies used by subjects in the Exit and the Stay treatments.

Our main hypothesis is in regard to a comparison of the efficiency levels observed in Exit treatments with those in Stay treatments. Figure 7 depicts the efficiency in each treatment in the second half of the experiment, 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%. Moreover, there is no statistical difference between the efficiency levels observed in the Stay 25 and Stay 30 treatments ($p > 0.10$).

The increase in efficiency levels in Games 25 and 30 explains the differences in final payoffs of players between Exit and Stay treatments reported in Table 9, 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 the Stay than in the Exit treatment. This result is driven by the difference in efficiency levels, as weak players are less likely to remain unmatched in the Stay than in the Exit treatment. However, conditional on reaching an efficient final match, strong players obtain higher shares, and weak players lower shares, in the Stay than the Exit treatment. This is also intuitive. The Stay treatment should make alternative matches more valuable, and it is the strong players who benefit from these alternatives.[33]

Figure 7: Efficiency of the final match, last five repetitions

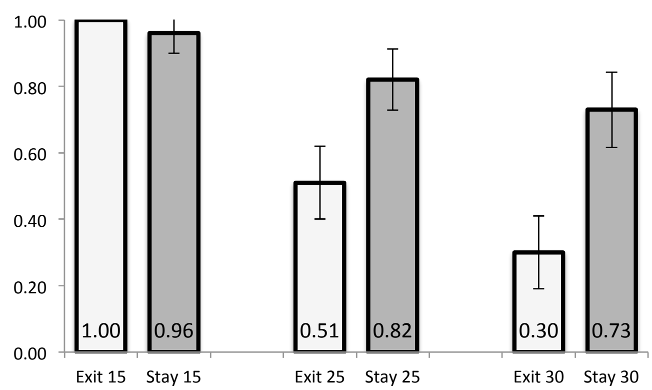

Notes: Average efficiency per treatment is reported, along with the 95% confidence interval, computed using robust standard errors, where errors are clustered at the session level. We focus only on groups that finished the game naturally and not because of the 1% termination probability.

We next ask whether the shift in efficiency levels documented in Figure 7 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.

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

Table 9: Payoffs of players by network position, last five repetitions

	Game 15		Game 25		Game 30	
	B (C)	A (D)	B (C)	A (D)	B (C)	A (D)
All Final Matches						
Exit 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						
Exit 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.

For each game, Figure 8 compares the CDFs of individual frequencies, in the Stay and Exit treatments, of proposing efficiently by strong players in the second half of the experiment.[34] 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.[35] Regression analysis confirms these results: $p < 0.01$ in both the Exit 25 vs. Stay 25 regression and the Exit 30 vs. Stay 30 regression, while $p > 0.10$ in Exit 15 vs. Stay 15 regression.

Figure 8: CDFs of frequency of efficient proposals by strong players, last five repetitions

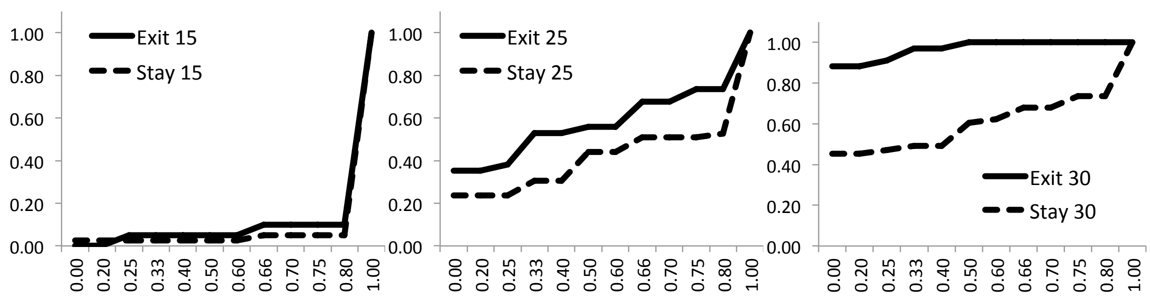

Notes: We present the cumulative distribution functions summarizing individual frequencies of proposing efficiently in the first round of the last five repetitions in each session when a subject performed a role of strong player. The horizontal axes indicate the likelihood of proposing efficiently, while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

Our theoretical conjecture described in Section 7 is based on the idea that the

[34]In Figure 8 we focus on the first-round behavior of strong players in each repetition in the second half of the experiment. Including all rounds in which all players were unmatched generates a figure very similar to Figure 8 (see the Supplementary Appendix, Section 2.2).

[35]For example, in the Exit 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.

possibility of renegotiation converts the short-lived outside options of players into long-lived ones that can be treated as exogenous, since market opportunities are never lost. For this reason we expect unmatched players' payoffs to be approximately independent of who is matched. To illustrate this idea, consider Game 30 and player A's payoff agreed upon in the match between A and C, depending on the circumstances in which this agreement was formed: One situation corresponds to the match between A and C being formed when B and D are both unmatched, and the other corresponds to the situation in which B and D have already formed a match of their own. If our conjecture is correct, then we expect player A's payoffs to be the same in both these situations. This is contrary to the dynamics we documented in the Exit treatments, in which strong players who reach an agreement first enjoy a premium over strong players who reach an agreement second. We test this conjecture by comparing payoffs of strong players in the Stay treatments in the two market compositions described above. Regression analysis confirms our hypothesis: In both the Stay 25 and Stay 30 treatments, we detect no significant difference, at the 5% level, between payoffs of strong players in the two situations described.[36]

We conclude this section by classifying groups that reached 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 last five repetitions in each session, 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 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 exploit profitable deviations and rematch with their efficient partners.[37] The behavior of these groups seems to us more consistent with the core theory than with our conjectured behavior, while the behavior of the other groups is in line with what we were expecting.

9 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 from the market and nevertheless find that the evolution of players' bargaining positions as others reach agreements is enough to drive significant mismatch.

[36]Note that these regressions include both amounts accepted by strong players when they were the responders and amounts received by strong players from the agreements they proposed themselves.

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

We turn off search frictions by letting players be patient, remove information problems by giving everyone symmetric information about match surpluses and remove coordination problems by considering very simple settings.

When players exit after agreements are reached we document extensive inefficiencies and show that the Markov perfect equilibria describe play well. We show that empirically the terms agreed depend on who is present in the market. Players' proposals of, and acceptances of, inefficient matches can then be rationalized by the concern they will find themselves in weaker bargaining positions in the future. Further investigating this mechanism we find that permitting players to renege on existing matches and rematch for a small cost reduces mismatch. This change holds the market context more stationary preventing players' bargaining positions from evolving as before. Allowing players to renege increases the frequency with which players make offers to and accept offers from their efficient partners.

To permit the rigorous application of non-cooperative game theory, we specify a bargaining protocol for our experiments. While non-cooperative theories abstract from dynamics altogether, this gives the non-cooperative theory an advantage over cooperative theories insofar as it is tailored to the experiment. Nevertheless, the key force driving inefficiencies in our experiment is the evolution of the market as players exit, and that feature appears to be present empirically. It would be interesting to see whether by placing less structure on the protocol the market is able to limit inefficiencies. A second line of inquiry that would be interesting, but again raises design challenges, would be to study the interaction of and relative magnitudes of different frictions. By making the discount factor lower search frictions can be introduced, by limiting the information available to the participants asymmetric information can be introduced and by considering larger markets coordination problems can be created.

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.

Baron, D. P. and J. A. Ferejohn (1989): "Bargaining in legislatures," *American political science review*, 83, 1181–1206.

Becker, G. S. (1973): "A theory of marriage: Part I," *Journal of Political Economy*, 813–846.

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

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

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.*

Cole, H., G. Mailath, and A. Postlewaite (2001a): "Efficient non-contractible investments in finite economies," *Advances in Theoretical Economics*, 1.

——— (2001b): "Efficient non-contractible investments in large economies," *Journal of Economic Theory*, 101, 333–373.

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.

Driessen, T. S. (1998): "A note on the inclusion of the kernel in the core of the bilateral assignment game," *International Journal of Game Theory*, 27, 301–303.

Elliott, M. (2014): "Heterogeneities and the fragility of the labor market," *mimeo.*

——— (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.*

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

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.

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.

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.

Shimer, R. and L. Smith (2000): "Assortative matching and search," *Econometrica*, 68, 343–369.

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

Ståhl, I. (1972): "Bargaining theory," *Economic Research Institute at the Stockholm School of Economics (EFI),*.

Sutton, J. (1986): "Non-cooperative bargaining theory: An introduction," *The Review of Economic Studies*, 53, 709–724.

A Derivation of Theoretical Predictions

A.1 Symmetric Pairwise Bargained Outcomes

In the four-player games we study, the system of equations we need to solve to find the fixed point reduces to

$$\begin{aligned}
u_A &= 0 + \frac{1}{2}(20 - 0 - (s_{BD} - u_B)) \\
u_B &= (s_{BD} - v_D) + \frac{1}{2}(20 - (s_{BD} - v_D) - 0) \\
v_C &= 0 + \frac{1}{2}(20 - 0 - (s_{BD} - v_D)) \\
v_D &= (s_{BD} - u_B) + \frac{1}{2}(20 - (s_{BD} - u_B) - 0)
\end{aligned}$$

for $s_{BD} \in \{15, 25, 30\}$. Solving this system of equations yields the predictions stated.

A.2 Markov Perfect Equilibria

In this section we derive the MPE that we test. A more formal and general derivation is provided in Nava and Elliott (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.

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

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

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

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 $20 - 4.17 < 30 - 13.33$.

A.3 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 any player in the unique perfect equilibrium of the subgame where all other players except their efficient partner has exited.

We construct an efficient perfect equilibrium in which, on path, player i offers their efficient partner $\mu^*(i)$ a payoff δV_i^M and their partner accepts. 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 their 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 them with a payoff of δV_S^M. As in the MPE the strong players mix between making offers to each other and offering to their efficient partner, the indifference condition implies that $25 - \delta V_S^M = 20 - \delta V_W^M$. Thus the 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 will 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 Game 30 the strong players strictly prefer offering to each other than offering efficiently in the MPE such that $30 - \delta V_D^M > 20 - \delta V_C^M$. 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

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

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

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

But then if selected as the proposer, A can either stick with the prescribed strategy 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 necessary because we can impose punishment only if the offer 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 9. 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

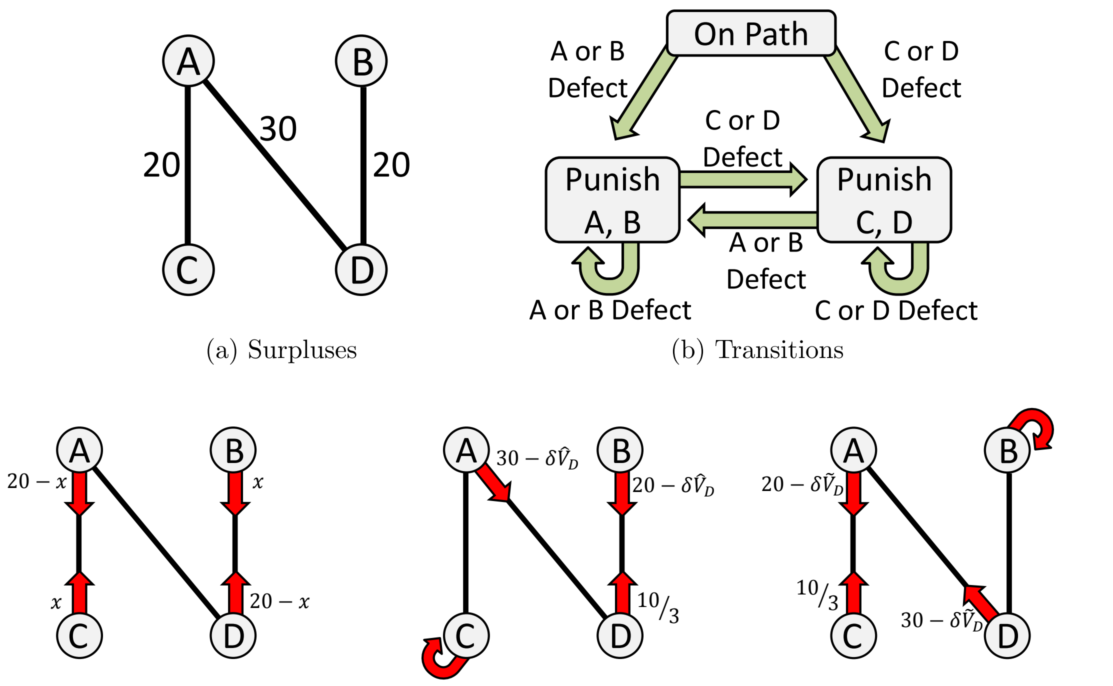

(a) Surpluses (b) Transitions

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

Figure 9: 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.

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, and unlike with MPE reversion, these transitions also reward all the players to whom the punished player is linked. These transitions are illustrated in panel (b) of Figure 9.

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 $13\frac{1}{3} \leq 13\frac{1}{3}$. This inequality becomes strict, while the others remain satisfied, when the offer D makes to B is increased slightly.[38] 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 $(10\frac{5}{9}, 13\frac{8}{9})$.

This construction of 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 will 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 $(10\frac{5}{9}, 12\frac{2}{9})$.

[38]We make this incentive constraint tight to find the full range of payoffs these punishment strategies can support.