Assignment Markets: Theory and Experiments

Arthur Dolgopolov, Daniel Houser, Cesar Martinelli and Thomas Stratmann

January 25, 2020

ASSIGNMENT MARKETS: THEORY AND EXPERIMENTS[1]

Arthur Dolgopolov[a], Daniel Houser[a], Cesar Martinelli[a] and Thomas Stratmann[a]

We study theoretically and experimentally assignment markets: two-sided markets where indivisible heterogeneous items with unit demand and unit supply are traded for money. These types of environments are important as they reflect key characteristics of, for example, housing markets. Shapley and Shubik (1971) proved that core allocations solve the optimal assignment problem and correspond one-to-one to competitive equilibrium allocations. Despite this result appearing a half-century ago, it remains unknown whether core allocations in assignment markets can be reached by decentralized individual behavior. This paper takes a step towards filling that gap. We explore, theoretically and experimentally, how different trading institutions help agents discover and reach optimal assignments in a noncooperative environment. In view of the complexity of the market, experimental results are perhaps surprisingly well-predicted by theory. We observe market outcomes close to Nash equilibrium predictions under auction-like institutions, and close to generalized bargaining for institutions that feature decentralized communication.

Keywords: assignment, housing market, bargaining, laboratory experiment.

1. INTRODUCTION

The assignment game Shapley and Shubik (1971) models markets characterized by indivisibility and heterogeneity of goods, and inflexibility of supply and demand satiated at exactly one unit. Allocations in this game consist of assignments of buyers to sellers and trade between matched buyers and sellers. This game not only provides general insights about exchange but captures key features of a number of important markets.

[1]The authors thank Jordi Brandts for helpful comments that improved the paper.

[a]ICES and Department of Economics, George Mason University

Arthur Dolgopolov: adolgopo@gmu.edu (corresponding author)

Daniel Houser: dhouser@gmu.edu

Cesar Martinelli: cmarti33@gmu.edu

Thomas Stratmann: tstratma@gmu.edu

For example, the assignment game well-portrays defining features of housing markets. To see this, note that a typical housing market includes multiple sellers, each of whom have one house they hope to trade for money. The market also includes multiple buyers, each of whom has demand for exactly one house, and is willing to exchange money for a house. Importantly, the same house is valued differently by different buyers, because each buyer places different importance on each house's characteristics (e.g., location, size). This heterogeneity can be amplified by informational asymmetries among buyers created by real-estate agents (or perhaps dampened by publicly available databases such as Zillow). Ultimately, the impact of the discrete supply and demand structure, valuation heterogeneity and information differences are modulated by the institution mediating trade, which in housing markets is typically bilateral bargaining.

Shapley and Shubik (1971) proved that core allocations solve the optimal assignment problem and correspond one-to-one to competitive equilibrium allocations. Despite this result appearing a half-century ago, it remains unknown whether core allocations can be reached by decentralized individual behavior in assignment markets. This paper is a step towards filling that gap. We explore, theoretically and experimentally, how different trading institutions help agents discover and reach optimal assignments in a noncooperative market environment.

Our theory benchmark is a strategic market game adapted from Dubey (1982), Simon (1984), and Benassy (1986). Buyers and sellers are allowed to bid simultaneously, with a clearing house picking an allocation consistent with the bid and ask prices submitted respectively by buyers and sellers. We show that every Nash equilibrium allocation is a competitive equilibrium allocation of an economy composed exclusively of the goods traded. That is, inefficiencies may arise due to coordination failures—buyers and sellers may fail to trade some goods that would be traded in an optimal assignment. In an alternative strategic game and following Pérez-Castrillo and Sotomayor (2017), we allow sellers to post prices before buyers are allowed to bid. Subgame perfection in this, now sequential, game selects the best competitive equilibrium for sellers, and therefore an optimal assignment. We also explore the generalized bargaining solution in our setting, and show that both core allocations and the leximin allocation (restricted to money transfers between paired traders) can be reached for different weights.

These three theoretical models naturally correspond to three market institutions, each with their own merits: competitive pressure in the markets with unstructured negotiation may not be as strong as as in the structured bidding, but provides more freedom for renegotiation. We test the performance of institutions and accuracy of predictions in each case in the laboratory. The experimental protocol is kept compatible with both bargaining and the core in all treatments, i.e. it is the same assignment problem implemented as different institutions.

For our experimental tests, we adapt the three-buyer, three-seller numerical example from the classic work by Shapley and Shubik (1971). In this example, the competitive equilibrium assignment and therefore the optimal assignment is unique; as is generally the case, equilibrium prices form a lattice from the best price vector to sellers to the best price vector for buyers. We implement three different treatments in the lab. In the first treatment, we adapt a double auction as introduced by Smith (1962) and now common in market experiments with homogeneous goods (see e.g. the survey by Friedman (2018)). In the second treatment, we allow sellers to post prices before buyers bid for their goods, as in the sequential strategic game described above. In the third treatment, we allow for bilateral, private, free format communication between buyers and sellers, mimicking real-estate interactions, with the possibility of striking deals that traders can renege from by making another deal. In the three treatments, we do not inform traders of the valuations of other traders, and we repeat the experiment several rounds with fixed roles and valuations, to allow for learning of market-relevant information and for experience.

We find that the double auction reaches the (unique) competitive assignment, i.e. the optimal assignment, 29% of the time. (As a comparison, the other five assignments allocating all three goods occur 3% of the time.) A second best assignment allocating only two goods, which is also a Nash outcome, occurs 37% of the time. Non-Nash assignments occur 22% of the time. The posted price treatment does slightly better, with the optimal assignment being reached 38% and the second best assignment being reached 29% of the time, and non-Nash assignments occurring 25% of the time. Finally, in the treatment with communication the optimal assignment is reached 59% of the time and the second best assignment is reached 6% of the time. Something remarkable is that the leximin assignment, which is not a Nash assignment, occurs 21% of the time in the treatment with communication, while it is fairly rare in the other treatments. In terms of splitting the gains from trade, in the posted price treatment sellers appropriate 64% of the surplus, versus 61% in the double auction treatment and 57% in the treatment with communication. In the treatment with communication, payoffs are close to the core when the optimal assignment is reached, and close to leximin payoffs (requiring equal split of gains from trade in matched pairs) when the leximin assignment is reached.

Summarizing, the posted price treatment favors the sellers, who are able to reap some gains from the ability to commit. However, in spite of being induced by the unique subgame perfect equilibrium, the optimal assignment is not reached in the posted price treatment as frequently as in the treatment with communication. As in other experiments,[1] communication favors coordination, leading to the optimal assignment, but also introduces either bargaining or pro-social behavior, in our case, exemplified by the frequency of the leximin matching.

2. LITERATURE REVIEW

Related literature can be organized into three blocks of decreasing abstraction: theory and general properties of the core in assignment games, strategic models, and experimental work. This section reviews all three, but the contribution of this paper lies with the second and third blocks—decentralized non-cooperative models and experimental comparison of their predictions without limiting the analysis to the core region.

Assignment problems have been studied using linear programming techniques starting with Von Neumann and Morgenstern's (1953) discussion of exchange, motivated

[1] See e.g. Martinelli and Palfrey (forthcoming).

in part by Böhm-Bawerk's (1891) study of a competitive market. Gale (1960) proves the existence of competitive equilibria in assignment problems. As mentioned above, Shapley and Shubik (1971) show that there is a one-to-one mapping from the core to competitive equilibria, which also possess a lattice structure that orders the allocations from buyer-best to seller-best. These polytopes, "45° lattices" in terms of Quint (1991) characterize exactly the set of all possible cores of assignment games. Ostroy (2018) shows the differences between the core and competitive equilibria in market games (and absence of this difference for assignment markets) through differentiability with respect to individuals and generators of positively homogeneous functions.

The allocations themselves can be easily obtained through linear programs because assignment problems have a totally unimodular constraint matrix. This ensures that though optimal assignments are always integer, they can be obtained through linear programming without integer constraints. A summary of results for assignment problems is offered by Roth and Sotomayor (1990, chapter 8) and more recently by Núñez and Rafels (2015). The close connection between unimodularity, linear programming and competitive equilibria has also been recently explored in Baldwin and Klemperer (2019).

While a powerful cooperative solution concept for the assignment problem with a non-empty set of predictions and an axiomatic characterization (see e.g. Toda (2005)), the core of the assignment game requires a particular market process to be implemented in a non-cooperative game. Theoretical work usually proceeds by making regularity assumptions to ensure that the core is exactly the whole set of predicted outcomes. Kamecke (1989) describes a version of non-cooperative characterization of an assignment market based on demand game from Nash (1953), but imposes an exogenous cost of making high unfulfilled offers in the game to avoid noncompetitive Nash equilibria. We move in the opposite direction to remain compatible with real-world markets where behavior may deviate from competitive outcomes and some markets may not open.

Consequently we limit the assumptions to natural technical restrictions for both non-cooperative solutions that we consider—the strategic market game and bargaining. In the case of the market game, the clearing house prefers more trade if there is no arbitrage (Simon, 1984). Indeed, a clearing house that actively precludes efficient trade for exogenous reasons is not in the spirit of the problem. For the bargaining solution, we do not require equal negotiation power, but allow it to vary between all players. An alternative interpretation is the varying probability of each player making an offer at each step in the corresponding structured bargaining game (Okada, 1996, 2010). Equipped with free communication, the players may not necessarily take strict turns making offers, and if some players are more active in the negotiation process, it naturally shifts the set of subgame-perfect equilibria. However, in line with our minimal assumptions approach, we do not introduce deviations by coalitions to the bargaining game as the papers above, keeping the core a separate concept.

From a strategic perspective, Schotter (1974) contains an early discussion of incentives for competitive behavior under different auctioning rules. Demange and Gale (1985) show that if a mechanism implements the optimal allocation for one side of the market, it may be manipulable by coalitions of players on the other side of the market. Pérez-Castrillo and Sotomayor (2002, 2017) provide sequential mechanisms where sellers announce their prices before buyers can bid for the goods, and show that such mechanisms implement the optimal allocation for sellers. In this paper we contribute to the strategic literature on assignment problems by studying the set of equilibria of market games. The equivalence result between (active trade) Nash and competitive equilibrium allocations obtained by Dubey (1982), Simon (1984), and Benassy (1986) does not hold in our setting because of market thinness, so characterizing the set of Nash equilibria is illuminating regarding possible outcomes of trading.[2] Additionally, we also consider predictions based on other solution concepts, namely leximin and Nash bargaining solution. Envy-free allocations in the assignment problem have been studied by Alkan et al. (1991) as a centralized problem under a budget constraint. Here, on the other hand, we consider a market situation without external subsidies.

[2] Market games in the context of indivisible, homogenous goods have been considered by Friedman and Ostroy (1995) and Martinelli et al. (2019).

From an experimental perspective, the heterogenous indivisible goods environment offers a larger spectrum of institutional possibilities than classic works on double auctions and pit trading markets, and it remains largely unexplored. While there is an extensive experimental literature for various market environments, studies of behavior in market settings with multiple goods and indivisibilities have largely been limited to centralized mechanisms, e.g. Rassenti et al. (1982). Bilateral negotiations similar to real-estate interactions are present only in a few studies. Experimental studies mimic this process either as communication over the telephone (Hong and Plott, 1982; Grether and Plot, 1984), using private booths (Crössmann, 1982; Selten, 1970), or with subjects walking around the room as in the original experiment by Chamberlin (1948) (see also Plott, 1982). Unlike Selten (1970), we do not allow private communication between sellers or between buyers to conspire against the other side of the market. We implement three new institutional set-ups that differ in sequence of moves, communication and market clearing procedures, motivated by housing and labor markets. The comparisons between our treatments are also relevant to the discussion of recontracting in Edgeworth and Walrasian *tâtonnement* found in Walker (1973), with some experimental evidence for positive effect of opportunities for recontracting on efficiency (Smith et al., 1982).

It is illustrative to compare the results to experimental studies of double auction and pit trading markets with homogeneous goods. The efficiency in the literature is shown in Figure 1. We did not include the experimental designs that combine the two mechanisms into a centralized market with an option to trade off-floor (Campbell et al., 1991) or endogenous market structures where traders decide who will be informed about their offer (Kirchsteiger et al., 2005, directed bid-ask market treatment).

There is a common belief that double auctions outperform decentralized bargaining in terms of total surplus achieved by traders. The evidence for this here is mixed, but it is also suggested by within-study comparisons, e.g. in Kirchsteiger et al. (2005). Interestingly, our assignment markets point to the contrary—since renegotiation is now a valuable instrument for efficiency, pit trading treatment outperforms centralized

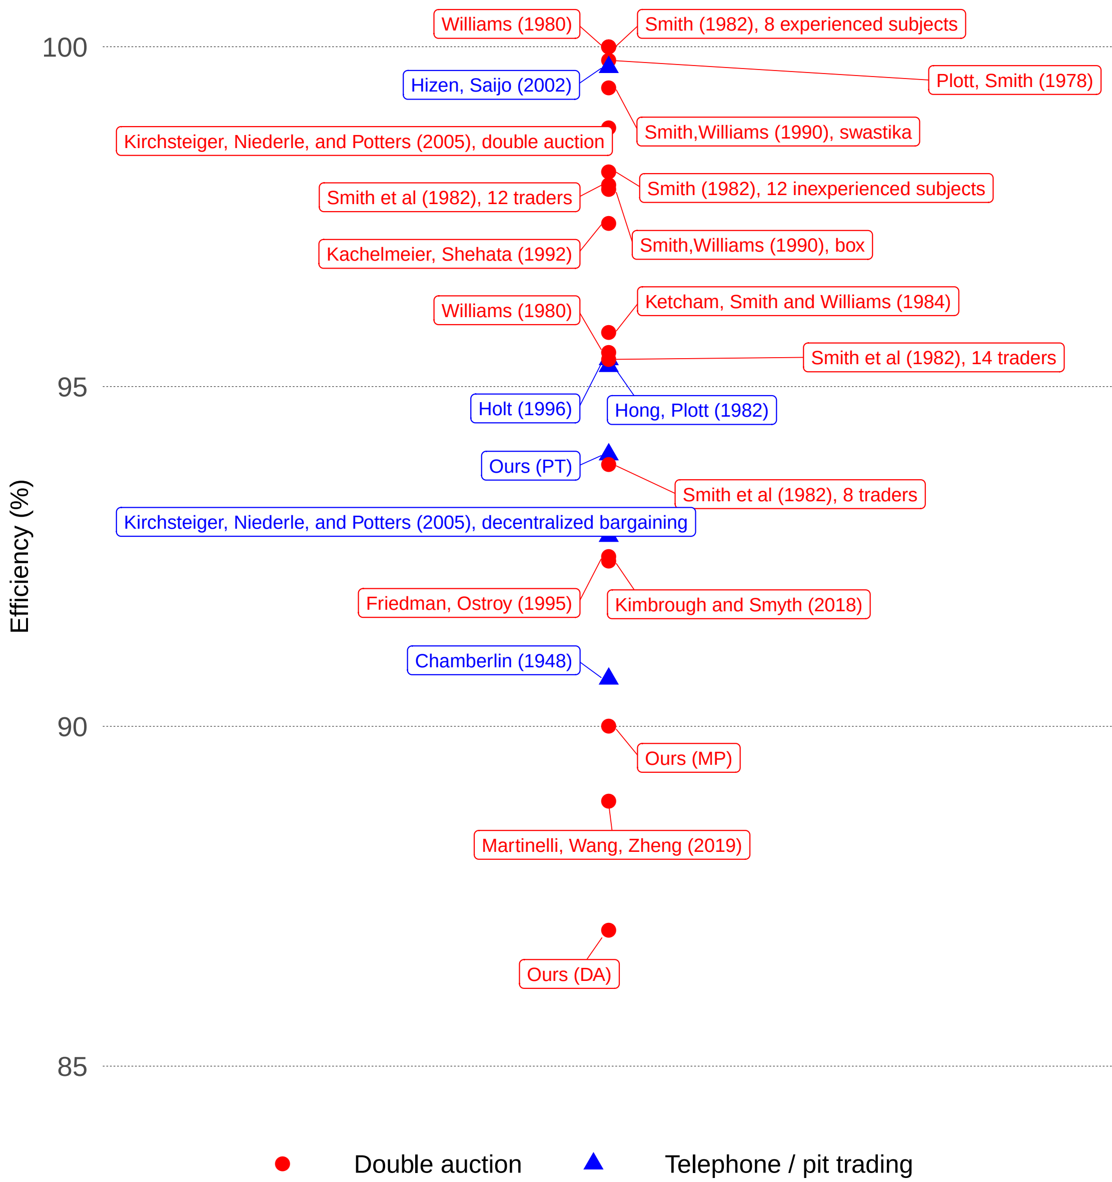

Figure 1: Efficiency in the literature on double auctions and pit trading (telephone) negotiations

market treatments. Moreover this is achieved, in line with previous bargaining experiments, through higher volume at more equalized prices. Of course this is illustrative at best since there is no reasonable way to aggregate efficiency across studies because of the differences in trading protocols, gains from trade, the number of traders, and complexity of finding the efficient matching.

The remainder of the paper is organized as follows. Section 3 is split into three parts. We first introduce the assignment game and review the core and competitive equilibrium results from Shapley and Shubik (1971). We then develop the non-cooperative counterpart, the strategic market game, and finally discuss generalized Nash bargaining and leximin solutions. Section 4 describes the experimental design and motivates the three treatments within the housing market setting. Section 5 relates theoretical predictions to experimental results. Section 6 concludes.

3. THEORY AND PREDICTIONS

3.1. *The assignment game*

Let B and S denote respectively the set of buyers and sellers, with $|B| = N$ and $|S| = M$. Each seller has one good to sell, and therefore we use S for the set of goods as well. In accordance with the experimental framing, we refer to the goods as 'widgets.' Each buyer j has a valuation for each seller's widget i, denoted by h_{ij}, while each seller has a reservation value for his own widget c_i. Possible market operations are transfers of widgets from sellers to buyers, and transfers of money from one agent to another. Each buyer j has a valuation for each seller's widget i, denoted by h_{ij}, while each seller has a reservation value for his own widget c_i. Payoffs are linear in money; in particular, if buyer j buys a widget from seller i and pays p_i, the buyer obtains payoff $h_{ij} - p_i$ and the seller obtains payoff $p_i - c_i$. We can refer to (S, B) as an economy.

Following Shapley and Shubik (1971), we define the characteristic function $v(\cdot)$ of the assignment cooperative game to be the smallest superadditive function on the power set of $B \cup S$ satisfying $v(C) = 0$ if $C \subseteq B$ or $C \subseteq S$ and $v(\{i, j\}) = \max(0, h_{ij} - c_i) \equiv a_{ij}$ if $i \in S$ and $j \in B$. The interpretation of $v(\cdot)$ is that the best a coalition can do is assigning buyers to sellers within the coalition, and splitting the gains from trade. An *assignment* for the whole market can be represented by an $M \times N$ binary matrix x with $x_{ij} = 1$ whenever good i is allocated to buyer j, and $x_{ij} = 0$ for all j if the good i is not sold. The value of the grand coalition $v(B \cup S)$ is obtained by solving the optimal assignment problem, that is

$$\begin{gathered} \max \sum_{i\in S} \sum_{j \in B} x_{ij} a_{ij}, \text{ subject to} \\ \sum_{j\in B} x_{ij} \leq 1 \text{ for all } i \in S \quad \text{and} \quad \sum_{i \in S} x_{ij} \leq 1 \text{ for all } j \in B. \end{gathered} \tag{3.1}$$

Using standard arguments, the maximum is attained and the constraint $x_{ij} \in \{0, 1\}$ is not binding at the maximum.[3]

A *core imputation* is a payoff vector $(\hat{u}, \hat{v}) = ((\hat{u}_i)_{i\in S}, (\hat{v}_j)_{j\in B})$ that is feasible for the grand coalition, that is $\sum_{i\in S} \hat{u}_i + \sum_{j\in B} \hat{v}_j \leq v(B \cup S)$, and such that no coalition $C \subseteq B \cup S$ can benefit all its members by deviating and splitting among them the value $v(C)$ of the coalition.

Shapley and Shubik (1971) show that the set of core imputations coincide with the set of solutions to the dual program to problem 3.1:

$$\begin{gathered} \min \sum_{i\in S} \hat{u}_i + \sum_{j \in B} \hat{v}_j, \text{ such that} \\ \hat{u}_i + \hat{v}_j \geq a_{ij} \text{ for all } (i, j) \in S \times B, \text{ and } \hat{u}_i \geq 0, \hat{v}_j \geq 0. \end{gathered} \tag{3.2}$$

The constraints $\hat{u}_i + \hat{v}_j \geq a_{ij}$ in problem 3.2 are binding whenever the corresponding buyer and seller are matched in the optimal assignment, that is when $x_{ij} = 1$ in a solution to program 3.1. Intuitively, payoffs of every core imputation can be achieved by monetary transfers from a buyer to a seller who are optimally assigned to each other.

Throughout we will use an example borrowed from Shapley and Shubik (1971), which is described in Table I. Every row in the table is a seller, and every column is a

[3]As pointed by Shapley and Shubik (1971), the optimal assignment is generally unique, as in the example below. There could be several optimal assignments in special cases, for instance if several widgets are perfect substitutes for buyers.

TABLE I

Widget values for buyers and sellers

Widgets (j)	Seller's reservation value (E$) ($c_j$)	Buyers' valuations (E$) ($h_{1j}$)	(h_{2j})	(h_{3j})
1	360	460	520	400
2	300	440	480	420
3	380	420	440	340

buyer. For this example, the values a_{ij} comprise the following matrix A, where each element is the joint maximal payoff of buyer j and seller i in experimental dollars (E$):

$$
\begin{array}{cc} & \text{(buyers)} \\ \text{A=} \ \text{(sellers)} & \begin{bmatrix} 100 & 160 & 40 \\ 140 & 180 & 120 \\ 40 & 60 & 0 \end{bmatrix} \end{array} .
$$

The unique optimal assignment is obtained by solving problem 3.1, and is given by $x_{12} = x_{23} = x_{31} = 1$ and $x_{ij} = 0$ otherwise. Optimal matches are shown circled in matrix A. We will use a shorter notation below, denoting assignments by three digit numbers. Each digit is the number of the good the buyer bought, e.g. the optimal assignment in the example can be written as [*312*], where buyer 1 bought widget 3, buyer 2 bought widget 1 and buyer 3 bought widget 2. If a buyer does not buy anything, we will write zero.

Fixing the optimal assignment, we can find core imputations by using the constraints in problem 3.2. Since the constraints corresponding to optimal matches are binding, we have that all core imputations satisfy $u_1 + v_2 = 160$, $u_2 + v_3 = 120$, and $u_3 + v_1 = 40$. Projecting all core imputations into u space, we obtain the pentahedron depicted in Figure 2a. In particular, the buyer-optimal imputation is $(u_*, v^*) = ((60, 100, 0), (40, 100, 20))$ and the seller-optimal imputation is $(u^*, v_*) = ((100, 120, 20), (20, 60, 0))$. In between these two extremes lies the core of the game,

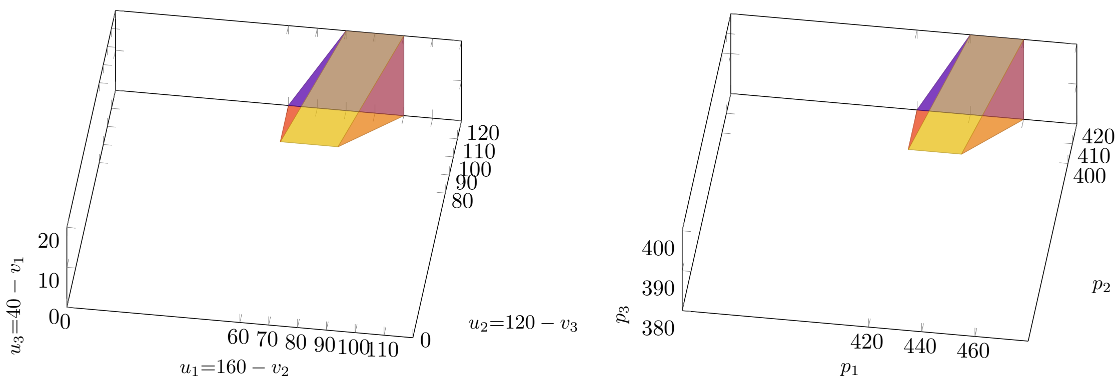

a) Core imputations in sellers' payoffs space

b) Competitive equilibrium price vectors

Figure 2: Core and competitive equilibria

which has a convenient structure. Let the partial order $(u, v) \leq_S (u', v')$ capture seller-optimality, i.e. $(u, v) \leq_S (u', v') \iff u_i \leq u'_i$ for all $i \in S$. Combined with this partial order, the core is a complete lattice.

To describe competitive behavior, let $Y_j = \{0, 1\}^M$, with typical element $y_j = (y_{ij})$, represent the set of possible demand vectors for buyer $j \in B$, with the interpretation that $y_{ij} = 1$ if buyer j demand widget i and $y_{ij} = 1$ otherwise. That is, a buyer can acquire one widget, several, or none. Similarly, let $Y_i = \{0, 1\}$, with typical element y_i, be the set of possible supply decisions by seller i, with the interpretation that $y_i = 1$ if i sells her widget and $y_i = 0$ otherwise.

A *competitive equilibrium* for the economy (S, B) is a pair (y, p), where $y = ((y_i)_{i\in S}, (y_j)_{j\in B}) \in \prod_{i\in S} Y_i \times \prod_{j\in B} Y_j$ and $p = (p_i)_{i\in S} \in \Re^M_+$ such that

$$\max_{i\in S}(y_{ij}h_{ij}) - \sum_{i\in S} p_i y_{ij} \geq \max_{i\in S}(y'_{ij}h_{ij}) - \sum_{i\in S} p_i y'_{ij} \text{ for all } y'_j \in Y_j, \text{ for all } j \in B,$$

$$(p_i - c_i)y_i \geq (p_i - c_i)y'_i \text{ for all } y'_i \in Y_i, \text{ for all } i \in S, \text{ and} \tag{3.3}$$

$$\sum_{j\in B} y_{ij} = y_i \text{ for all } i \in S.$$

The first set of conditions represent utility maximization by buyers, and encode the assumption that buyers can enjoy at most one widget. The second set of conditions represent profit maximization by sellers, and the third set of conditions are market clearing conditions for each of the widgets.

Shapley and Shubik (1971) show that in any assignment game, including the previous example, all core imputations $(\hat{u}, \hat{v})$ can be supported in competitive equilibria (y, p) by prices

(3.4) $\quad p_i = \hat{u}_i + c_i$ for all $i \in S$.

That is, a core imputation and the corresponding competitive equilibrium implement the same (optimal) assignment, and the money transfers implicit in the core imputation are implemented by market prices. Conversely, all competitive equilibrium payoffs are core imputations. We therefore have a direct one-to-one mapping from core imputations to equilibrium price vectors,[4] as illustrated by a similar pentahedron in Figure 2b.

3.2. *Strategic market game*

We can now define the strategic solution. In the strategic *market game* each seller $i \in S$ submits a price $p_i^S \in \mathbb{R}_+$, and each buyer submits an M-vector of positive bids $p_j^B \in \mathbb{R}_+^M, p_j^B = (p_{1j}^B, ..., p_{Mj}^B)$. The sets of admissible prices/bids for player $k \in B \cup S$ is denoted W_k. An *offer profile* combines actions of all players in a tuple $w = (p_1^B, \ldots, p_N^B, p_1^S, \ldots, p_M^S) \in W = \prod_{k \in B \cup S} W_k$.

Once all bids and prices are submitted, a clearing house chooses an assignment (or a lottery over assignments) in the feasible set

$$X = \{(x_{11}, \ldots, x_{NM}) : x_{ij} \in \{0, 1\}, \sum_{j \in B} x_{ij} \leq 1 \text{ for all } i \in S\}.$$

The clearing house allocates trade to maximize surplus $\pi(x, w) = \sum_{i \in S} \sum_{j \in B} x_{ij}(p_{ij}^B - p_i^S)$, i.e. it draws from the following set of surplus-maximizing assignments:

$$Y(w) = \{x \in X : \pi(x, w) \geq \pi(x', w) \text{ for all } x' \in X\}.$$

Once the clearing house chooses an assignment x, the market clears at the buyers'

[4]The one-to-one relation described by equation 3.4 holds for widgets that are optimally assigned; for widgets that are not assigned to any buyer, any price $p_i \geq c_i$ is competitive.

prices,[5] and buyers and sellers get the payoffs

$$\max_{i\in S}(x_{ij}h_{ij}) - \sum_{i\in S} p^B_{ij}x_{ij} \quad \text{and} \quad (p^B_{ij} - c_i)\sum_{j\in B} x_{ij},$$

respectively. To ensure that the clearing house prefers more trade even when arbitrage is zero, we assume that it chooses assignments that are not ray-dominated (Simon, 1984):

$$F(w) = \{x \in Y(w) : \text{there is no } \phi \in Y(w) \text{ such that}$$
$$\phi \neq x \text{ and } \phi_{ij} \geq x_{ij} \text{ for all } i \in S, j \in B\}.$$

We also assume that the clearing house randomizes over the full $F(w)$, that is it chooses randomly according to some distribution that has positive probability on all $F(w)$.

To describe the set of Nash equilibria of the market game, note that for every subset $S' \subseteq S$ of sellers (including the empty set), we can define an economy (S', B) with S' as the set of sellers and B as the set of buyers. Let (y', p'), where

$$y' = ((y'_i)_{i\in S'}, (y'_j)_{j\in B}) \in Y_{S',B} \equiv \prod_{i\in S'} Y_i \times \prod_{j\in B} Y_j \quad \text{and} \quad p' = (p'_i)_{i\in S'} \in \Re^{|S'|}_+$$

be a competitive equilibrium for any such economy. With a slight abuse of notation, for any $y \in Y_{S',B}$, let

$$x(y) \equiv (x \in X : x_{ij} = y'_{ij} \text{ if } i \in S' \text{ and } j \in B, \text{ and } x_{i'j} = 0 \text{ if } i' \notin S' \text{ and } j \in B).$$

Note that if $x(y') \in F(w)$ for some bid profile $w = (p^B_1, \ldots, p^B_N, p^S_1, \ldots, p^S_M)$ such that

$$y'_{i'j} = 1 \quad \Rightarrow \quad p^S_i = p^B_{ij} = p'_i \quad \text{for every } i' \in S' \text{ and } j \in B,$$

then $x(y')$ induces the same allocation than the competitive equilibrium (y', p'); that is, it induces the same assignment and the same money transfers than (y', p') for every $i \in S'$ and $j \in B$, with the remaining sellers $i' \notin S'$ remaining unassigned.

[5] Allocating surplus to one side of the market follows the convention in Dubey (1982) and also reflects the experimental treatment where one side of the market has the full market power.

We say that a bid profile w induces the same allocation than (y', p') if $F(w)$ is a singleton satisfying $F(w) = \{x(y')\}$, and moreover, for every $i' \in S'$ and $j \in B$ such that $y'_{i'j} = 1$, we have $p_i^S = p_{ij}^B = p'_i$.

We have

THEOREM 3.1 *If (y', p') is a competitive equilibrium for some economy (S', B) for some $S' \subseteq S$, then there is a Nash equilibrium bid profile w that induces the same allocation.*

PROOF: To prove the theorem, consider w such that

$$p_i^S = \begin{cases} p'_i & \text{if } i \in S' \\ \kappa & \text{if } i \notin S' \end{cases}$$

for some $\kappa > \max_{i \in S, j \in B} h_{ij}$, and

$$p_{ij}^B = \begin{cases} p'_i & \text{if } y'_{ij} = 1 \\ 0 & \text{otherwise} \end{cases}.$$

While the assignment $x(y')$ makes zero surplus, every other assignment makes zero or negative surplus, and moreover the assignment above ray-dominates every other zero-surplus assignment. Hence, $F(w) = \{x(y')\}$, as desired.

It is straightforward to check that, since ask prices for widgets such that $i \in S'$ are competitive, and prices for widgets $i \notin S'$ are prohibitively expensive, no buyer has an incentive to deviate from w to any another bid profile inducing an assignment different from x with positive probability. Similarly, since bid prices for widgets such that $i \in S'$ are competitive, and prices for widgets $i \notin S'$ are zero, no seller has an incentive to deviate from w to any another bid profile inducing an assignment different from x with positive probability. *Q.E.D.*

An immediate corollary of Theorem 3.1 is that every competitive equilibrium allocation of the original economy (and hence any corresponding core imputation) can be supported by a Nash equilibrium. However, other allocations can be supported as well; in fact, recalling that in our environment every widget represents a different market, any arbitrary subset of markets can be shut down in a Nash equilibrium.

This is the result of a coordination failure—intuitively, markets for particular widgets may not open because each side of the market expects the other side not to show up. Such allocations are the semi-Walrasian allocations in terms of Mas-Colell (1982). Mas-Colell (1982) shows that these allocations are more difficult to destabilize than other non-equilibrium allocations. That is if some active market is not an equilibrium, the allocation can be blocked by one type of traders ("1-blocking"), but if all active markets are in equilibrium, then at least m types of traders may be needed (m-blocking), with m no more than and sometimes exactly the number of inactive markets.

A converse result to theorem 3.1 also holds.[6] We have

THEOREM 3.2 *If $F(w) = \{x\}$ and w is a Nash equilibrium, then there is a competitive equilibrium (y', p') for some economy (S', B), where $S' \subseteq S$, such that x induces the same allocation than (y', p').*

PROOF: Suppose that $F(w) = \{x\}$ and $w = (p_1^B, ..., p_N^B, p_1^S, ..., p_M^S)$ is a Nash equilibrium. First, we claim that surplus is driven down to zero, i.e.

$$\sum_{i \in S} \sum_{j \in B} x_{ij}(p_{ij}^B - p_i^S) = 0 \text{ for all } x \in F(w).$$

f If, on the contrary, for some $j \in B$ and $i \in S$, $x_{ij} = 1$ and $p_{ij}^B > p_i^S$, then buyer j has a profitable deviation to $p_{ij}'^B = p_{ij}^B + \epsilon$ for small enough $\epsilon > 0$ so that x is still preferable for the clearing house after the deviation to any $x' \neq F(w)$. Hence, $x_{ij} = 1$ implies $p_i^S = p_{ij}^B$, and every possible match makes zero or negative surplus.

Now let S' be the subset of sellers whose goods are assigned by x, let y' be the solution to $x(y') = x$, and let $p' = (p_i^S)_{i \in S'}$. We claim that (y', p') is a competitive equilibrium for the economy (S', B). To see this, note that market clearing is guaranteed by the definition of $F(w)$. Profit maximization for each seller i at the given

[6]Theorem 3.2 can be extended to Nash equilibria with random outcomes as long as the clearing house cannot react to an arbitrarily small price reduction by a buyer from an initial strategy profile w by increasing the probability of a different assignment which has the same surplus in the initial situation and excludes the buyer.

price p_i^S is guaranteed by the fact that since w is a Nash equilibrium, $p_i^S \geq c_i$ so that selling at the price p_i^S is at least as good as not selling. Finally, we claim that each buyer $j \in B$ maximizes utility by choosing y'_j given prices p'. For suppose not; then one of the following cases must hold:

(i) $x_{ij} = 1$ for some $i \in S$ but there is $i' \in S$ such that $h_{i'j} - p_{i'}^S > h_{ij} - p_i^S$,

(ii) $x_{ij} = 0$ for all $i \in S$ but there is $i' \in S$ such that $h_{i'j} - p_{i'}^S > 0$,

(iii) $x_{ij} = 1$ for some $i \in S$ but $h_{ij} - p_i^S < 0$.

In case (i), buyer j can deviate to $p'^B_{i'j} = p_{i'}^S + \epsilon$ and $p'^B_{i''j} = 0$ for all $i'' \neq i$ for

$$0 < \epsilon < (h_{i'j} - p_{i'}^S) - (h_{ij} - p_i^S).$$

After the deviation, the clearing house should match j and i' since every other match makes zero or negative surplus by the previous step. Similarly, in case (ii), buyer j can deviate to $p'^B_{i'j} = p_{i'}^S + \epsilon$ and $p'^B_{i''j} = 0$ for all $i'' \neq i$ for

$$0 < \epsilon < h_{i'j} - p_{i'}^S.$$

As in the previous case, the clearing house should match j and i' after the deviation. In case (iii), buyer j can benefit by deviating to $p'^B_{i'j} = 0$ for all i', which guarantees a payoff of zero. Hence, in each of the three cases, w cannot be a Nash equilibrium. *Q.E.D.*

For economies like our running example in Table I, in which the optimal assignment assigns all sellers, the following corollary follows from previous results. Intuitively, if all markets open, the strategic game leads to a competitive equilibrium for the complete economy.

COROLLARY 1 *If w is a Nash equilibrium such that $F(w) = \{x\}$ satisfying $\sum_{j \in B} x_{ij} = 1$ for all $i \in S$, then x is a solution to problem 3.1.*

Consider again our running example. From the previous corollary, the only complete assignment that can be induced (with probability one) by a Nash equilibrium is the optimal assignment [*312*]. Deleting rows in matrix A, we get the other assignments that can be induced deterministically by Nash equilibria, ordered by total payoff gains: [*210*], [*320*], [*310*], [*020*], [*010*], [*030*], and [*000*]. The assignment [*210*], shown circled in matrix A below,

$$A = \begin{bmatrix} 100 & 160 & 40 \\ 140 & 180 & 120 \\ 40 & 60 & 0 \end{bmatrix},$$

is the closest to the optimal assignment in total payoff among all suboptimal assignments. It can be supported, for instance, by the following Nash equilibrium profile:

$$\begin{aligned} w &= (p_1^S, p_2^S, p_3^S, p_1^B, p_2^B, p_3^B) \\ &= (470, 430, 440, (460, 430, 380), (470, 430, 380), (400, 420, 380)). \end{aligned}$$

In other words, these assignments are semi-Walrasian equilibria by Theorem 3.2. The bipartite nature of the assignment market simplifies the analysis considerably compared to Mas-Colell (1982)—two players at the most need to change actions simultaneously to destabilize any Nash equilibrium. However, the transition from fixing the coordination problem in a single pair may affect many or all the agents—while we only need to change actions by the third buyer or the third seller and their partners in the assignment [*210*] to nudge the markets toward the competitive equilibrium, the matching among active traders changes as well—buyer 1 is rematched with seller 3 instead of seller 2. In fact, [*210*] is the second most frequent assignment, after the optimal assignment, in the Double Auction and Minimum Price experiments in the following sections.

As an alternative noncooperative game, inspired by Pérez-Castrillo and Sotomayor (2002, 2017), consider a sequential game with complete information in which sellers are allowed to choose simultaneously their prices first, and then buyers choose simultaneously their bids, with the clearing house choosing the final allocation as before. By standard arguments, every Nash equilibrium of the simultaneous game corresponds to a a Nash equilibrium of the sequential game in which the buyers choose the same bid no matter what happens in the first stage of the game. More interestingly, following Pérez-Castrillo and Sotomayor, there is a unique subgame perfect equilibrium path and it leads to the best allocation for sellers in the core, corresponding to the imputation (u^*, v_*).

3.3. *Bargaining*

Bargaining models are a candidate for predicting behavior especially for treatments with communication. The set of feasible payoff imputations under the restriction to bilateral transfers among matched pairs can be written as $U \in \Re_+^{M+N}$, with typical element $(u, v) = (u_1, \ldots, u_M, v_1, \ldots, v_N)$, such that there exists $x \in X$ such that $u_i + v_j = h_{ij} - c_i$ if $x_{ij} = 1$, $u_i = 0$ if $\sum_{j \in B} x_{ij} = 0$, and $v_j = 0$ if $\sum_{i \in S} x_{ij} = 0$. The generalized Nash bargaining solution[7] normally leads to the following convex optimization problem:

$$\max_{(u,v) \in U} \prod_{i \in S} u_i^{\alpha_i} \prod_{j \in B} v_j^{\alpha_j}, \tag{3.5}$$

for some weights $\alpha_i, \alpha_j \geq 0$ such that $\sum_{i \in S} \alpha_i + \sum_{j \in B} \alpha_j = 1$. We will use the solutions to this optimization program as our bargaining solution concept, which coincides with the classical Nash bargaining solution for convex problems. Since core payoffs are in U, and maximize the sum of payoffs, it follows that for every imputation in the core there are weights such that the imputation solves problem 3.5, but there are solutions for other weights which are not in the core.

[7]The non-convexity of the bargaining problem may seem disconcerting to the reader because the classic axiomatic Nash bargaining solution is not well-defined and there are several competing extensions. We follow a direct extension of the Nash Bargaining solution similar to Kaneko (1980) and Herrero (1989). That is, we fall back to a set-valued solution concept that returns maximizers to the weighted Nash function, possibly under Pareto optimality constraint. Alternatively, if we follow the approach in Conley and Wilkie (1996) the Pareto and Nash bargaining sets in Figure 3 will coincide. Finally, a third option is to add lotteries to the bargaining set, which may not be satisfactory if agents do not have expected utility preferences. The difference is contained in the non-convex regions with imputations that are Pareto dominated by lotteries—whether these imputations are reasonable outcomes depends on the chosen axioms. Since we are interested in the relationship between the bargaining solution, the core and Nash equilibria, we omit the axiomatic details relating the bargaining solution to the Pareto set. In other words, Theorem 3.3 is robust to other, less restrictive, definitions.

TABLE II

Generalized Nash bargaining with random weights

Assignment	Frequency
[*312*] (Efficient)	58.8%
[*132*] (3rd best)	35.1%
[*231*]	3.6%
[*321*]	2.5%

Note: Based on 1,000 draws with weights drawn from a uniform distribution.

Although we refer to assignments other than the core assignment as suboptimal, since they do not maximize the sum of utilities in a transferable utility environment, note that every outcome of generalized Nash bargaining is by definition Pareto efficient, since varying the weights assigned to the traders traces the utility-possibility frontier under the restriction to bilateral transfers among matched pairs.

In our running example, the solution with equal weights for all traders corresponds to the optimal assignment, but leads to the following price vector (transfers to sellers from matched buyers) $p = (440, 360, 400)$, which is not competitive. if we draw weights randomly, other (suboptimal) complete assignments will have positive probability, while incomplete assignments will have zero probability as long as weights are positive. Drawing weights randomly from a uniform distribution, the most likely outcome is the efficient core assignment with 60% probability, and then the assignment [*132*] with 37% (see Table II). We refer to [*132*] as a third best since it is follows in total payoff after [*312*] and [*320*] (and tied with [*120*]).

While the Nash equilibria correspond exactly to the semi-Walrasian equilibria in any assignment market, the relationship between Nash bargaining solution, the core and Pareto optima is less straightforward, and is summarized in Figure 3. In particular, any core imputation can be supported by a generalized Nash bargaining solution, but not all Pareto optima generally have this property. In fact, the Pareto set is not convex, unless the problem is trivial, and all Pareto optima are obtainable from the same assignment.

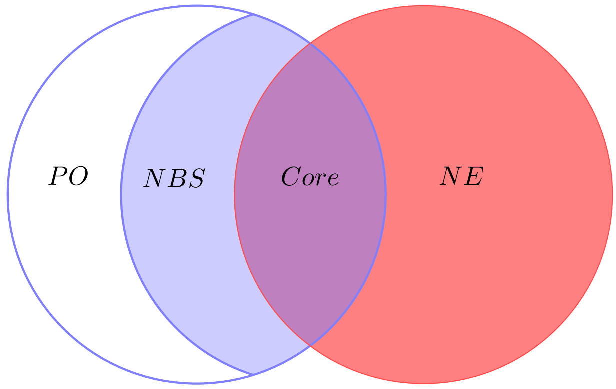

Figure 3: Relationship between the Pareto optima (PO), generalized Nash bargaining solution (NBS), the core, and Nash equilibria (NE)

THEOREM 3.3 *Every core imputation is a generalized Nash Bargaining solution for some vector of weights α.*

PROOF: Suppose (u^*, v^*) is some core imputation, attainable at some core assignment x. Let $\hat{U}$ be the maximized surplus obtained in the core. Consider the generalized Nash bargaining problem of sharing $\hat{U}$ between all agents with weight vector a, $a_i = \frac{u_i^*}{\hat{U}}, a_j = \frac{u_j^*}{\hat{U}}$, for $i \in S$, $j \in B$, disregarding the roles and market structure. This is a standard convex optimization problem over a compact set:

$$\max_{(u,v):\sum_{i \in S} u_i + \sum_{j \in B} v_j \leq \hat{U}} \prod_{i \in S} u_i^{\alpha_i} \prod_{j \in B} v_j^{\alpha_j}, \tag{3.6}$$

Solutions to this problem are well-defined and can be obtained from first order conditions: $u_i = a_i \hat{U} = u_i^*$ and $v_j = a_j \hat{U} = v_j^*$ for all $i \in S, j \in B$. Since $U \subseteq \{(u, v) : \sum_{i \in S} u_i + \sum_{j \in B} v_j \leq \hat{U}\}$, the imputation (u^*, v^*) also maximizes the Nash product over a smaller set U. Therefore (u^*, v^*) is the generalized Nash bargaining solution for weights a. *Q.E.D.*

This logic does not extend to other Pareto optima that do not maximize the total surplus. To see that the set of generalized Nash bargaining solutions does not span the whole Pareto set, consider an assignment market with two buyers and two sellers, represented by the following matrix:

$$A = \begin{bmatrix} 2 & 1 \\ 3 & 4 \end{bmatrix}$$

Consider the imputation $(u, v) = ((0.5, 0.6), (2.4, 0.5))$. This outcome is attainable in the circled matching, and it is also Pareto optimal. To see this, note that this outcome is Pareto optimal for the circled matching (utilities sum up to circled elements), and buyer 1 is getting a utility level of 2.4, which is infeasible in the other full matching. It is however straightforward to check that it can not be supported by any generalized Nash bargaining solution.

Finally, let us consider the leximin solution defined over the set U. Leximin may be an attractive criterion for egalitarian traders; see e.g. Sen (1970). Leximin requires maximizing lexicographically the worst-off trader, then the second worst-off, etc. We can find the leximin solution by splitting surplus equally in every buyer-seller pair for every complete assignment, and then maximizing iteratively the worst off matched pair of traders and fixing their utility for the consecutive optimization until we get the payoffs for all traders. Again, by choosing appropriate weights, we can obtain the leximin solution as a solution to problem 3.5. In our example, leximin leads to the third-best assignment [*132*], and delivers a payoff of 30 to the two worst-off traders.

The leximin solution is by definition Pareto optimal; however, it does not have to be a generalized Nash bargaining solution or be in the core. For example, for the assignment market below, the imputation $(u, v) : ((1, 1), (1, 1))$ is attainable in the circled matching, and it is also the leximin.

$$A = \begin{bmatrix} 1 & 2 \\ 2 & 100 \end{bmatrix}$$

To see this, note that there is no way to have the worst-off agent obtain a payoff larger than 1 in this or any other matching. It is however again straightforward to check that it can not be supported by any generalized Nash bargaining solution, and it is not in the core.

4. EXPERIMENTAL DESIGN

4.1. *Treatments*

In our experimental treatments, we inform traders of the number of other buyers and sellers in the group, as well as each trader's own parameters (e.g. the vector $(h_{ik})_{i\in S}$ if trader k is a buyer, and c_k if trader k is a seller), but as in realistic market conditions we do not inform them about the parameters of other traders. Thus, behavior under the different treatments gives an indication about the ability of different market institutions to allow traders to discover enough information to induce competitive allocations, as those described in section 3.1, or to get stuck on some of the suboptimal Nash allocations described in section 3.2, representing the failure of one or several markets to open. The treatments are to some extent dynamic, and for each treatment we repeat the experiment with the same group and demand and supply conditions for fifteen rounds. Both features are there to facilitate learning; we can consider experience in previous rounds as providing knowledge about market conditions, in the way as for instance online databases with recorded past transactions. We consider three treatments, as detailed below.

4.2. *Double Auction (DA)*

The first treatment adapts the double auction commonly used in market experiments. In our version, the game is played in two stages. First, each seller sets a minimum price for his widget. This price has to be above his reservation value. When all minimum prices are set, the game proceeds to the trading stage. During the trading stage, buyers are allowed to bid for the widgets, and sellers are allowed to reduce their minimum prices. A buyer has a winning bid for a widget if the bid is the highest and it is above the current minimum price of the widget. To enforce unit demands, buyers who currently hold the winning bid for some widget are not permitted to make other bids until some other buyer outbids them. The round ends after 50 seconds of inactivity and buyers with winning bids will obtain their corresponding widgets. The game continues for fifteen rounds with groups, buyer and seller identities and reservation values unchanged between rounds, and with sellers revising their minimum prices at the first stage of each round. One round is then chosen at random for payment for each subject. The earnings for buyers are calculated as the difference between widget's value and the bid. Likewise, the earnings for the sellers equal the difference between the bid and the reservation value. Therefore sellers earn zero if they do not sell the widget and no trader is risking trading at a loss. Complete experimental instructions are available in the on-line appendix.

4.3. *Minimum Price (MP)*

The second treatment is set up to give market power to sellers, which should also help coordination. This treatment is based on Pérez-Castrillo and Sotomayor (2002, 2017) and relates to job market matching models in Crawford and Knoer (1981) and Kelso, Jr. and Crawford (1982). It also mimics commonly observed posted prices for houses with subsequent negotiation phase. As in the sequential game of section 3.2, the commitment ability of sellers should shift the results toward seller-optimal core allocations (upper-right corners in both panels in Figure 2). Like in the DA treatment, in the first stage each seller sets a minimum price for his widget. This price has to be above his reservation value, and this is the only way in which sellers actively participate in the market. When all minimum prices are set, the game proceeds to the trading stage with buyers bidding for the widgets.

4.4. *Pit Trading (PT)*

The third treatment is played through open-form bilateral communication between buyers and sellers. Every buyer has three chatboxes for private communication, one for each seller, and, similarly, each seller has three chatboxes for private communication, one for each buyer. Chatboxes allow exchanging messages and negotiating the price. (Temporary) deals are marked in the chatbox shared by a buyer and a seller. Buyers and sellers can communicate in the three chatboxes simultaneously and back out of a deal at any moment, possibly striking another deal. The current negotiated deals are finalized when the round ends. Traders in the same side of the market can only communicate with the other side of the market and can not communicate between themselves. This treatment is motivated by housing markets in which negotiations take place bilaterally over the phone or email. In this scenario, contracts are easier to back out from than in DA or MP.

Buyers' valuations and sellers' reservation values are the same in all treatments as in examples in the previous section (Table I), and ultimately lead to surpluses of coalitions in the matrix A. These payoffs follow the example in Shapley and Shubik (1971), but are scaled by a factor of 20 to form a large discrete space of integer bids and prices. The particular set-up in this table is attractive for the apparent complexity of the problem for the players, which can be seen through the simulations below.

4.5. *Simulations*

We emulate market dynamics in the DA treatment by making subjects make random or boundedly rational decisions in random order, while maintaining all the restrictions of the experiment (e.g. not losing money, only increasing bids for buyers or reducing prices for sellers, etc.). Since buyers with a standing bid cannot bid for a different widget until some other buyer outbids them, this process may lead to miscoordination even if buyers pick the good with the highest payoff and increment bids by the minimum amount. In particular, any complete assignment that does not result in a negative utility is possible. The zero in the third column of matrix A allows the third buyer and the third seller to be left unmatched in this process; note that the third buyer and the third seller have the lowest matching gains.

If subjects take their decisions in random order, with sellers doing small price increments and buyers doing random bid increases (zero-intelligence), the probability of arriving to the efficient allocation is only about 20% (see Table III). It is higher if buyers are boundedly rational and take optimal decisions given the current situation in the market (about 29%), and much higher only if buyers use only small bid increments. Arriving to competitive equilibrium prices by randomizing is extremely unlikely, except in the last case.

To provide an intuition for the high probability of the inefficient assignment [*210*],

TABLE III

SIMULATIONS

Assignment	Uniform	Highest margin	Highest margin + min. step
[*312*] (Efficient)	20.09%	28.76%	82.03%
[*210*] (Unique 2nd best)	19.86%	24.99%	2.04%
[*120*] (3rd best)	19.68%	21.84%	3.94%
[*321*]	15.23%	12.38%	3.35%
[*132*] (3rd best and leximin)	13.70%	6.92%	7.40%
[*231*]	11.44%	5.11%	1.25%
Competitive prices	0.08%	0.88%	40.15%
Competitive prices w/o good 3	3.16%	4.76%	0.51%
Efficiency	89.44%	91.07%	97.41%
Seller's share of surplus	72.86%	73.93%	50.80%

Note: Based on 50000 simulated markets. Players make moves in random order. Buyers randomly pick a good if they can earn a positive payoff by buying it, sellers gradually reduce the minimum price. Uniform: buyer's bid increase is drawn from uniform distribution over all values that give positive payoff. Highest margin: buyers bid for goods that give them the highest payoff at current prices, but randomize their bids uniformly. Highest margin + min. step: same, but buyers only increment bids by minimum amount.

note that a matching between buyer 1 and seller 2 can not happen in a competitive equilibrium, because if $u_2 \leq 120$ and $v_1 \geq 20$ then seller 2 would want to sell to buyer 3 instead, and if $u_2 \geq 100$ and $v_1 \leq 40$ then buyer 1 would want to buy widget 3 instead. Under the rules of the Double Auction (as well as under the rules of the Minimum Price treatment), if current prices are such that we are in the first situation, buyer 3 will be able to outbid buyer 2 and the market would move toward the equilibrium. However, if current prices are such that we are in the second situation, buyer 2 can not bid for good 3 since this buyer is already the highest bidder for good 2. Thus, the inability to renegotiate a deal may lead to inefficient assignments in the first two treatments.

4.6. *Summary of predictions*

Our predictions are motivated by the similarities between the theoretical models and the corresponding experimental treatments. Generally we expect the DA and MP treatments to be closer to the predictions of the structured models—the competitive equilibrium model and the core. On the contrary, the PT treatment, due to its open communication nature, is expected to be closer to unstructured bargaining model with predictions around focal points like leximin or other assignments compatible with bargaining solution in Table II. At the same time, we should see less miscoordination and unassigned goods than in DA and MP, which would in this case generate zero utility and be inconsistent with the Nash bargaining solution. Miscoordination is more likely in DA and MP, however, due to limited opportunities for communication and renegotiation.

Additionally, these predictions about the resulting assignments also restrict the possible areas where we would expect to see the realized prices. The price vectors are expected to be in the competitive equilibrium region in DA and MP treatments, with the latter likely closer to seller-optimal point, while the prices in the PT treatment are not restricted by the linear inequalities of the core region, and can lie further away from it. The predictions about prices are stronger hypotheses, nested within predictions about the assignments, and ultimately more powerful as indicated by power analysis in previous section.

More specifically, in the DA treatment, we expect core allocations to occur in the lab, as well as perhaps other allocations corresponding to Nash equilibria of the market game of section 3.2, with some markets possibly failing to operate.[8] Frequent occurrence of the third-best assignment [*120*] or other non-Nash assignments like [*321*] in addition to the two Nash assignments [*312*] and [*210*] would be inconsistent with the main theory, and would indicate that traders are using some heuristic strategies, locking themselves in suboptimal deals.

In the MP treatment, we expect core allocations close to the seller-optimal allocation, implying that the optimal assignment [*312*] should occur in the lab, as suggested by the subgame perfect equilibrium of the sequential game in section 3.2, and that payoffs for sellers should be higher than in other treatments. However, the sequential game has other Nash equilibria, in particular some inducing the second-best assignment [*210*].

Finally, we expect the opportunities for bilateral communication under the PT treatment to favor coordination in core allocations, implying that the optimal assignment [*312*] should occur in the lab more often than in other treatments. Alternatively, communication may lead to other allocations predicted by bargaining models, like [*132*]. Leximin also leads to [*132*], with a particular transfer structure.

5. EXPERIMENTAL RESULTS

Our experimental sessions were conducted with 180 GMU undergraduate subjects in total, broken down into groups of six (three buyers and three sellers) for a total of thirty groups. Experiments were conducted in oTree (Chen et al., 2016) from September, 2018 to April, 2019. The first 5 rounds were discarded as learning rounds, but this was not announced. The software matched participants in groups of 3 buyers and 3 sellers with the matching and roles unchanged over the entire experiment. Subjects were recruited for 90 minutes. Earnings were calculated in the experimental currency

[8] The DA treatment is a complex dynamic game in real time, so we take the strategic market game as a useful simplification.

TABLE IV

Efficiency across treatments (realized fraction of total surplus)

	Min	Max	Mean	Std. deviation (p.p.)
DA	0.00	1.00	0.87	0.18
MP	0.50	1.00	0.90	0.13
PT	0.63	1.00	0.94	0.09

TABLE V

P-values for rank-sum tests of efficiency between treatments

	DA	MP	Remaining obs.
DA			0.0113**
MP	0.2597		0.5599
PT	0.0011**	0.0282**	0.0018***

and converted to U.S. dollars at the end of the experiment at a rate of 5 E$ to 1 US$. Only the earnings from one randomly chosen round were paid, with a mean payment of $13.86 across treatments.

Table IV illustrates the mean, lowest, and highest payoff (as fraction of the maximum) for the three treatments. In this and other tables as well as in the statistical analysis below we drop the first five rounds as nosy learning. Consistent with the predictions above, the PT treatment performs better than the other two with respect to efficiency. The rank-sum statistical outlined in Table V confirm this: PT advantage over the other two treatments is statistically significative. Efficiency under the PT treatment (94%) is also above that under random behavior (89%) and random behavior with highest margin (91%), from Table III, although below the crawling algorithm described by the last column of Table III (97.5%).

Similarly, Table VII illustrates the the mean, lowest, and highest payoff to sellers as a share of realized payoff. Again, consistent with the predictions above, the sellers' share is largest under the MP treatment, and the rank-sum statistical outlined in Table VI confirm this. The MP treatment gives sellers an average share of 64% and the DA treatment an average share of 61%, equidistant from the average of the share under the seller-optimal (75%) and the buyer-optimal core imputation (50%). The

PT treatment is closer to equal split gains.

TABLE VI

P-values for rank-sum tests of seller surplus between treatments

	Min	Max	Mean	Std. deviation (p.p.)
DA	0.16	0.82	0.61	0.10
MP	0.35	0.91	0.64	0.09
PT	0.31	0.82	0.56	0.11

TABLE VII

Share of surplus that goes to the sellers

	DA	MP	Remaining obs.
DA			0.1859
MP	0.1142		$<0.0001^{***}$
PT	0.0001^{***}	$<0.0001^{***}$	$<0.0001^{***}$

Table IX shows frequency of different assignments under the three treatments, as well as the frequency of competitive equilibrium prices. We consider first the DA treatment. Note that the two top Nash assignments, [*312*] and [*210*], occur very frequently, adding up to two thirds of the time, as opposed to only about 40% of the time under random (uniform) simulated behavior (Table III), while the third-best allocation [*120*] whose frequency under random behavior is 20%, is rare in the lab. Looking at prices, while competitive prices are very rare under simulated behavior, they occur about a quarter of the time. As opposed to under random behavior, and as consistent with Nash equilibrium, the frequency of the optimal assignment is related to that of competitive prices.

Turning to the MP treatment, again the the top two Nash assignments add up to two thirds of the time. The efficient assignment is more frequent than under the DA treatment, but not as much to what subgame perfection would lead us to expect. The positive impact on efficiency is not large by the fact that the top two Nash assignments are close in total payoffs. (We compare prices under MP with other treatments below.)

Finally, turning to the PT treatment, we can see that the efficient assignment is twice as frequent as under the DA treatment, and considerably more frequent than under the MP treatment as well. This is the source of the efficiency advantage of the PT treatment. Note that the second best [*210*] happens with low probability as opposed to the other treatments, indicating that the ability to communicate makes the sort of coordination failure implied by that assignment less likely. What is remarkable is that the leximin and third best allocation [*210*], which is very unlikely under the other treatments, is the second most frequent under PT. This is an indication of either bargaining forces or prosocial behavior. This is consistent with our expectations given the nature of communication in the PT treatment. The low frequency of competitive prices is also striking. Another observation suggesting the bargaining nature of PT treatment is that prices arrived to by PT subjects do not group in regions where constraints bind as the DA and MP subjects do (see smaller panels in Figure 5 in the appendix). The high frequency of assignment [*210*] confirms the mis-coordination theory - this assignment is precisely a semi-Walrasian equilibrium in terms of Mas-Colell (1982) and a Nash equilibrium in Theorem 3.2 with third market closed.

The random simulations also allow us to calculate the Predictive Success Index (PSI, Selten (1991)), a measure that adjusts the frequency of successful predictions by the power of the test (the area of the core region in the present case). The values of PSI for the three treatments can be obtained by subtracting values for simulated power tests in Table III from actual results for human subjects in Table IX. Predictive success for the core price region and for the core matching are presented in Table VIII. Once again, competitive prices yield good predictions for DA and MP treatments. Consistent with expectations, for PT treatment a less restrictive prediction of having efficient matchings performs better when power is taken into account. That is full efficiency is reached more often, but not necessarily through the competitive price mechanism.

The three panels in Figure 4 present the detailed price dynamics for each of the treatments for each of the three widgets. Note that the first five rounds, before the dashed line in each panel, are not used for statistical purposes. Considering the three markets separately, prices tend to concentrate in the core region, marked in green, with some bias to prices below the core for widgets 1 and 2 and above the core for

TABLE VIII

Predictive Success

Theory	DA	MP	PT
Efficient matching	0.09	0.18	0.39
Competitive prices	0.24	0.16	0.1
Competitive prices for at least goods 1 and 2	0.37	0.3	0.09

TABLE IX

Frequencies of assignments across treatments (rounds 6–15)

Assignment	DA	MP	PT
[*312*] (Efficient; Nash)	29	38	59
[*210*] (Unique 2^{nd} best; Nash)	37	29	6
[*132*] (3^{rd} best and leximin)	2	4	21
[*120*] (3^{rd} best)	3	8	5
[*012*] (3^{rd} best)	8	6	4
[*310*] (Nash)	9	7	2
[*010*]	1	3	0
[*102*]	2	2	3
[*320*] (Nash)	2	1	0
[*321*]	1	1	0
[*130*]	0	1	0
[*200*]	4	0	0
[*032*]	1	0	0
[*000*] (Nash)	1	0	0
Competitive prices	24	16	10
Competitive prices w/o good 3	16	17	2
Total	100	100	100

widget 3. This is more remarkable in the case of the PT treatment given the frequency of the optimal assignment.

To further illustrate the price distribution under the different treatments, we calculate the average (minimum) Euclidian distance from observed price distributions to the core price region, conditional on the optimal assignment, the leximin assignment, and all assignments in Table X. Note that in general observed prices are closer to competitive prices under the DA and MP treatments, both conditioning on reaching

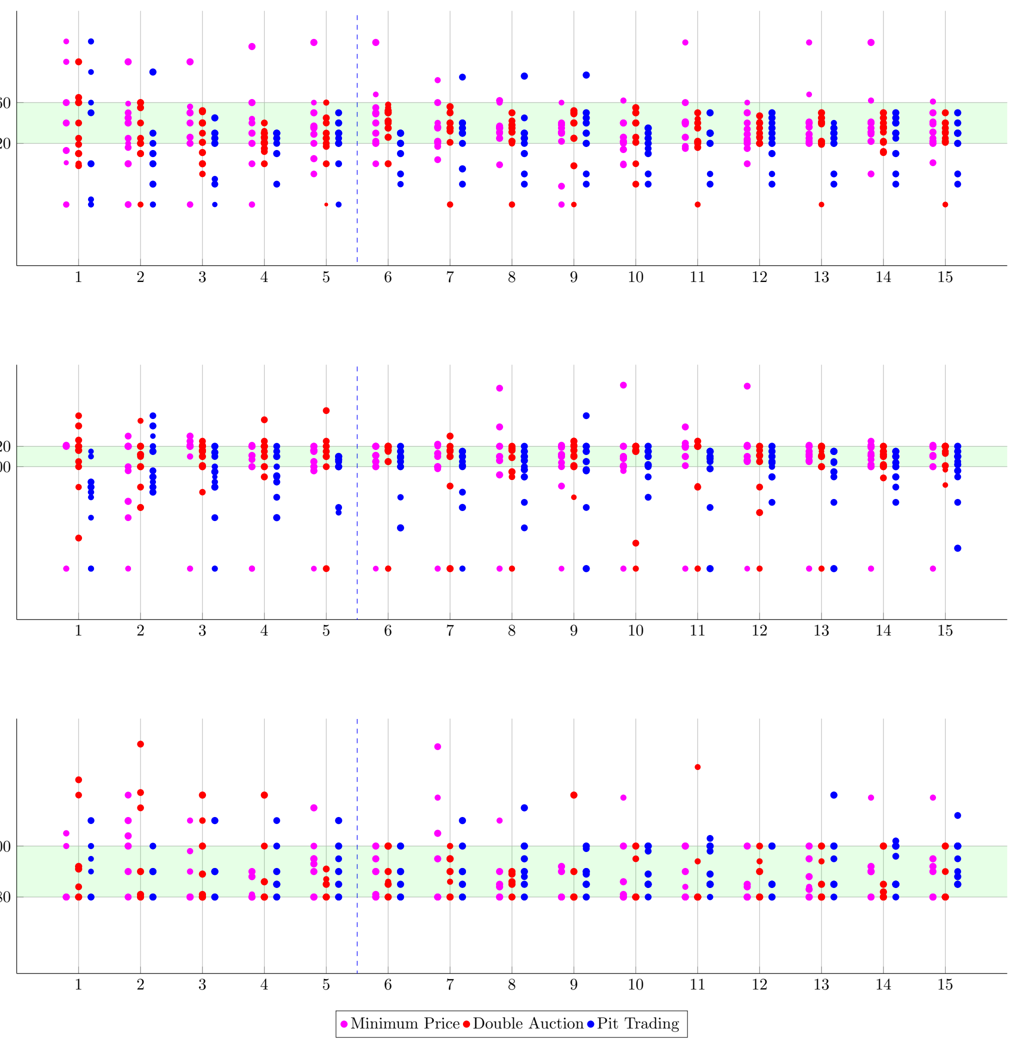

Figure 4: Resulting prices across rounds

Light green band is the range of core prices.

TABLE X

Average (minimum) distance between observed/simulated data and competitive prices. Standard deviation in parenthesis.

	Observed data			Simulations		
Assignment	DA	MP	PT	Uniform	Highest margin	Highest margin + min. step
Optimal	6.33	4.69	15.98	48.27	37.01	12.54
	(19.25)	(13.35)	(25.26)	(20.27)	(21.87)	(31.08)
Leximin	3.64	15.35	34.67	46.22	46.50	65.11
	(0.15)	(27.47)	(20.71)	(19.51)	(20.06)	(18.45)
All	20.93	19.06	22.09	183.61	208.59	43.44
	(35.62)	(36.32)	(26.76)	(162.44)	(168.95)	(96.52)

the optimal assignment and in general. Conditional on the leximin assignment, observed prices are much further away under the PT treatment; values conditional on the leximin assignment for the DA and MP treatment are included for completeness though the leximin assignment is fairly rare for those two treatments. Finally, within each treatment, the distance to competitive prices is much smaller conditional on reaching the optimal assignment than in general (the difference is statistically significant for the MP and the PT treatments). The distances in experimental data are also generally much smaller than in the simulations indicating again that the core is a good predictor and behavior can not be explained by random choice.

6. CONCLUSION

The inability of the core concept to capture the bargaining and strategic tactics was noted both in Shapley and Shubik (1971) and Von Neumann and Morgenstern (1953) in a critique of the classic Böhm-Bawerk (1891) solution for the horse market assignment problem. As noted by Shapley and Shubik (1971), the choice of the appropriate solution concept is dictated by the institutional form, including the communication structure. This logic is evident in the experiment.

Our experimental results notably follow the 70 years of intuition in experimental markets with negotiation—as first noted by Chamberlin (1948), the volume of trade in such "imperfect" markets is above the equilibrium level, and the prices are below the equilibrium level. Our experiment suggests that these effects extend to heterogeneous goods. We also observe a larger volume of trade and lower seller surplus when comparing the pit trading treatment with the other treatments. The higher volume of trade is consistent with the explanations based on imperfections in the matching procedure, which in our case also helps avoid the miscoordination of the semi-Walrasian outcomes. The usual explanations for the lower prices are behavioral—the familiarity of the buyer mentality or fairness of seller's prices as discussed by Chamberlin (1948) and in the subsequent literature. We complement these explanations with a bargaining model in the heterogeneous good setting that has a similar intuition and fits the data.

Assignment markets pose a difficult allocation problem, due to the heterogeneity of valuations for goods among the buyers, the fact that demand decisions are related, the thinness of markets, and the likely paucity of information. We illustrate some of these difficulties with a simplified strategic market model, and show that miscoordination outcomes, in which some markets fail to open, are possible in Nash equilibrium. We explore different market institutions in the lab, and show that under auction-like trading rules Nash outcomes are frequent. In particular, under double auction-like rules, about 78% of the time the assignment is consistent with Nash predictions, with the efficient assignment is reached 29% of the time. Giving the sellers some commitment power by allowing them to post prices before bidding by buyers increases the probability of the efficient assignment to 38%, as well as increases the share of the surplus going to sellers. Last, we consider a treatment that allows for bilateral, private communication between buyers and sellers, as well as provisional contracts that traders can back from, as in a real-estate market. In this last treatment, the efficient assignment is reached 59% of the time. Other assignments, which are consistent with generalized Nash bargaining but not with equilibria of the strategic game, become likely as well. A higher efficiency in the last treatment could be taken as a suggestive explanation for prevalence of open non-centralized negotiations in real estate, with the proviso that the opportunities for communication lead to trades at prices that differ from competitive prices, a possibility that indeed was implicitly entertained by both Von Neumann and Morgenstern (1953, p. 564) and Shapley and Shubik (1971, p. 128).

REFERENCES

Alkan, A., G. Demange, and D. Gale (1991): "Fair allocation of indivisible goods and criteria of justice," *Econometrica*, 59, 1023–1039.

Baldwin, E. and P. Klemperer (2019): "Understanding preferences: "Demand types", and the existence of equilibrium with indivisibilities," *Econometrica*, 87, 867–932.

Benassy, J.-P. (1986): "On competitive market mechanisms," *Econometrica*, 54, 95–108.

Böhm-Bawerk, E. V. (1891): *The Positive Theory of Capital*, G.E. Stechert and Co., New York.

Campbell, J., S. LaMaster, V. L. Smith, and M. Van Boening (1991): "Off-floor trading, disintegration, and the bid-ask spread in experimental markets," *Journal of Business*, 495–522.

Chamberlin, E. H. (1948): "An experimental imperfect market," *Journal of Political Economy*, 56, 95–108.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Conley, J. P. and S. Wilkie (1996): "An extension of the Nash bargaining solution to nonconvex problems," *Games and Economic Behavior*, 13, 26–38.

Crawford, V. P. and E. M. Knoer (1981): "Job matching with heterogeneous firms and workers," *Econometrica*, 49, 437–450.

Crössmann, H. J. (1982): *Entscheidungsverhalten auf unvollkommenen Märkten*, Barudio & Hess, Frankfurt a.M.

Demange, G. and D. Gale (1985): "The strategy structure of two-sided matching markets," *Econometrica*, 55, 873–888.

Dubey, P. (1982): "Price-quantity strategic market games," *Econometrica*, 50, 111–126.

Friedman, D. (2018): "The double auction market institution: A survey," in *The Double Auction Market: Institutions, Theories, and Evidence*, ed. by D. Friedman and J. Rust, Routledge, 3–26.

Friedman, D. and J. Ostroy (1995): "Competitivity in auction markets: An experimental and theoretical investigation," *Economic Journal*, 105, 22–53.

Gale, D. (1960): *The Theory of Linear Economic Models*, McGraw-Hill, New York.

Grether, D. M. and C. R. Plot (1984): "The effects of market practices in oligopolistic markets: an experimental examination of the ethyl case," *Economic Inquiry*, 22, 479–507.

Herrero, M. J. (1989): "The Nash program: non-convex bargaining problems," *Journal of Economic Theory*, 49, 266–277.

Hong, J. T. and C. R. Plott (1982): "Rate filing policies for inland water transportation: An experimental approach," *The Bell Journal of Economics*, 13, 1–19.

Kamecke, U. (1989): "Non-cooperative matching games," *International Journal of Game Theory*, 18, 423–431.

KANEKO, M. (1980): "An extension of the Nash bargaining problem and the Nash social welfare function," *Theory and Decision*, 12, 135–148.

KELSO, JR., A. S. AND V. P. CRAWFORD (1982): "Job matching, coalition formation, and gross substitutes," *Econometrica*, 50, 1483–1504.

KIRCHSTEIGER, G., M. NIEDERLE, AND J. POTTERS (2005): "Endogenizing market institutions: An experimental approach," *European Economic Review*, 49, 1827–1853.

MARTINELLI, C. AND T. R. PALFREY (forthcoming): "Communication and Information in Games of Collective Decision," in *Handbook of Experimental Game Theory*, ed. by R. Croson, M. Capra, M. I. Rigdon, and T. Rosenblat, Edward Elgar Publishing.

MARTINELLI, C., J. WANG, AND W. ZHENG (2019): "Competition with Indivisibilities and Few Traders," ICES, George Mason University.

MAS-COLELL, A. (1982): "Perfect competition and the core," *Review of Economic Studies*, 49, 15–30.

NASH, J. (1953): "Two-person cooperative games," *Econometrica*, 128–140.

NÚÑEZ, M. AND C. RAFELS (2015): "A survey on assignment markets," *Journal of Dynamics and Games*, 2, 227–256.

OKADA, A. (1996): "A noncooperative coalitional bargaining game with random proposers," *Games and Economic Behavior*, 16, 97–108.

——— (2010): "The Nash bargaining solution in general n-person cooperative games," *Journal of Economic Theory*, 145, 2356–2379.

OSTROY, J. M. (2018): "Competitive pricing and the core: With reference to matching," *Games and Economic Behavior*, 108, 558–573.

PÉREZ-CASTRILLO, D. AND M. SOTOMAYOR (2002): "A simple selling and buying procedure," *Journal of Economic Theory*, 2, 461–474.

——— (2017): "The outcome of competitive equilibrium rules in buyer-seller markets when the agents play strategically," *Economic Theory*, 64, 99–119.

PLOTT, C. R. (1982): "Industrial organization theory and experimental economics," *Journal of Economic Literature*, 20, 1485–1527.

QUINT, T. (1991): "Characterization of cores of assignment games," *International Journal of Game Theory*, 19, 413–420.

RASSENTI, S. J., V. L. SMITH, AND R. L. BULFIN (1982): "A combinatorial auction mechanism for airport time slot allocation," *Bell Journal of Economics*, 13, 402–417.

ROTH, A. AND M. SOTOMAYOR (1990): *Two-sided matching. A study in game-theoretic modeling and analysis*, Cambridge University Press, Cambridge UK.

SCHOTTER, A. (1974): "Auctioning Böhm-Bawerk's horses," *International Journal of Game Theory*, 3, 195–215.

Selten, R. (1970): "Ein Marktexperiment," in *Beiträge zur experimentellen Wirtschaftsforschung. Band 2*, ed. by H. Sauermann, J. C.B. Mohr (Paul Siebeck), Tübingen, 33–98.

——— (1991): "Properties of a measure of predictive success," *Mathematical Social Sciences*, 21, 153–167.

Sen, A. (1970): *Collective Choice and Social Welfare*, Holden-Day, San Francisco.

Shapley, L. S. and M. Shubik (1971): "The assignment game I: The core," *International Journal of Game Theory*, 1, 111–130.

Simon, L. K. (1984): "Bertrand, the Cournot paradigm and the theory of perfect competition," *Review of Economic Studies*, 51, 209–230.

Smith, V. (1962): "An experimental study of competitive market behavior," *Journal of Political Economy*, 70, 111–137.

Smith, V. L., A. W. Williams, W. K. Bratton, and M. G. Vannoni (1982): "Competitive market institutions: Double auctions vs. sealed bid-offer auctions," *American Economic Review*, 72, 58–77.

Toda, M. (2005): "Axiomatization of the core of assignment games," *Games and Economic Behavior*, 53, 248–261.

Von Neumann, J. and O. Morgenstern (1953): *Theory of Games and Economic Behavior*, Princeton University Press, Princeton NJ.

Walker, D. A. (1973): "Edgeworth's theory of recontract," *Economic Journal*, 83, 138–149.

APPENDIX A: PRICES IN EXPERIMENTAL MARKETS

The three panels in Figure 5 summarize prices for all observations in each of the treatments. The main picture in each of the panels presents the observed price vectors in three dimensions, as well as the core area, shaded in the picture. The smaller pictures present the projection of observed price vectors and the core area in two dimensions, for each pair of markets. The planes limiting the core area in the main picture, as well as the lines limiting the core area in the (projected) smaller pictures, represent various constraints on competitive prices. Intuitively, these constrains make sure that pairs of traders that should not be optimally assigned do not have a profitable opportunity to trade with each other, and that each trader prefers to trade rather than not. Slanted lines in the smaller pictures indicate are related to possible deviations from the core by pairs of traders, while lines that are parallel to either axis are related to participation constraints for the traders. Blue (entire) dots indicate complete assignments, while red dots indicate partial assignments, with the blank space indicating widgets that were not traded.

Panel (a) in Figure 5 illustrates that observed price vectors for complete assignments under the DA treatment are either in or very close to the core area. As the smaller pictures highlight, observed price vectors are often close or above various constraints. This is clear in particular for the price of widget 3 in relation to the price of the other widgets. This seemingly indicates resistance of seller 3 to competitive forces pushing down the price of the corresponding widget. Intuitively, in the optimal assignment, when the price of widget 3 rises too much, buyer 1 would be tempted to deviate and acquire instead widget 1 or widget 2.

Panel (b) in Figure 5 illustrates that the observed distribution of price vectors under the MP treatment is similar to that under the DA treatment, with a couple of exceptions. Constraints involving the prices of widgets 2 and 3 seem to be binding.

Panel (c) in Figure 5 illustrates a distribution of price vectors under the PT treatment which differs notably from the other two treatments. There are quite a few observed price vectors for complete assignments which are not close to the core. As illustrated by the smaller pictures, observed prices tend to violate core constraints in two particular dimensions: in several observations, the price of widget 3 is too high, and the price of widget 2 is too low. This is consistent with some of the observed complete market observations corresponding to the third best assignment [*132*]. In particular, leximin solution prescribes this assignment with the price vector $(460, 360, 410)$, which violates core constraints precisely because the price of widget 3 is too high, and the price of widget 2 is too low, respect to competitive prices.

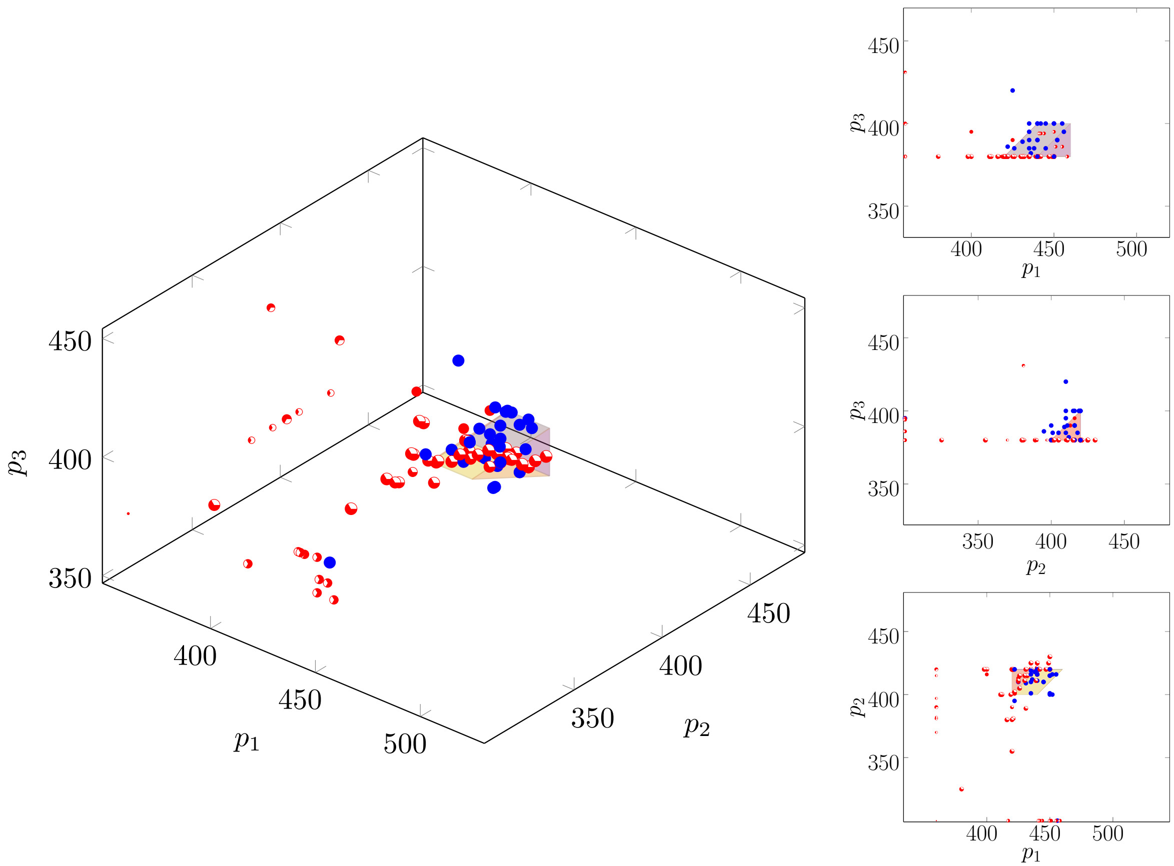

(a) Double Auction

Figure 5: Resulting prices across treatments

Note: blue points are rounds with maximized total utility, filling of the markers corresponds to items traded, size of the markers is the total utility. When item is not traded, the minimum price is used instead of the trading price.

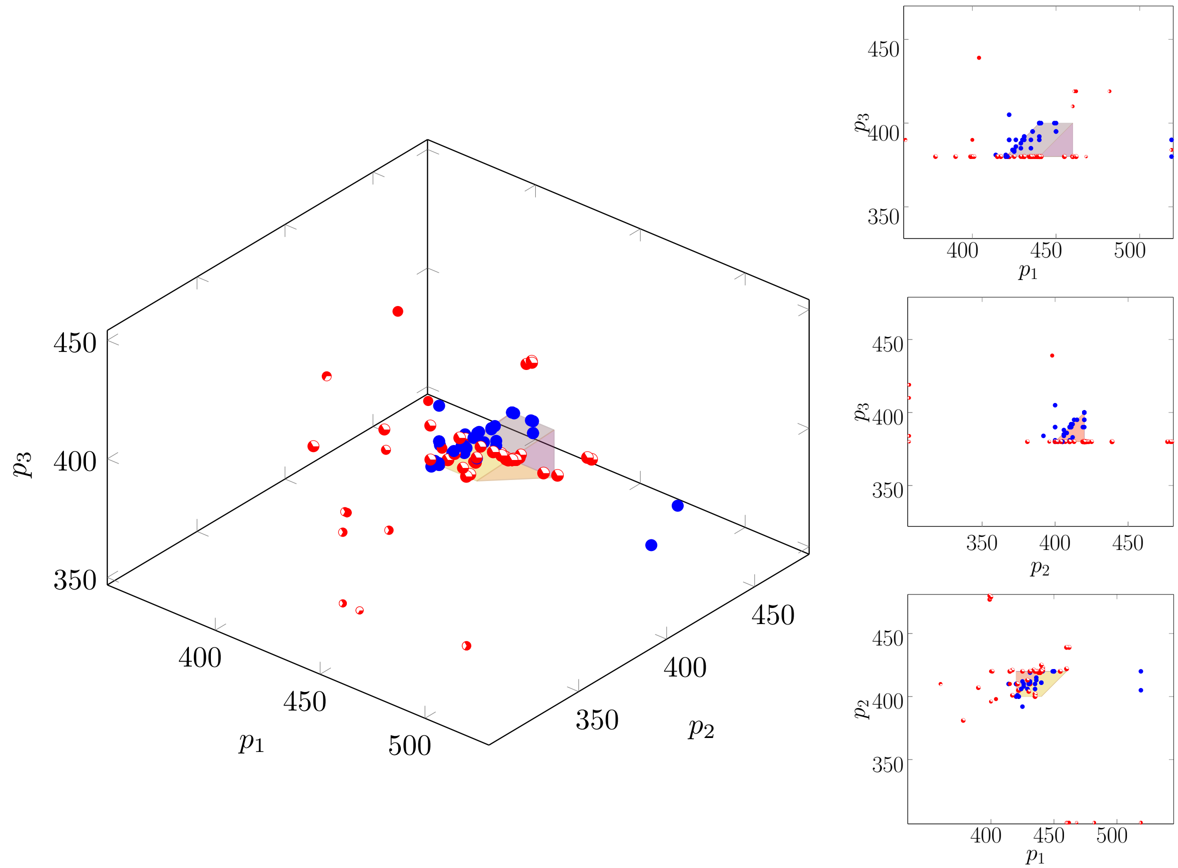

(b) Minimum Price

Figure 5: Resulting prices across treatments

Note: blue points are rounds with maximized total utility, filling of the markers corresponds to items traded, size of the markers is the total utility. When item is not traded, the minimum price is used instead of the trading price.

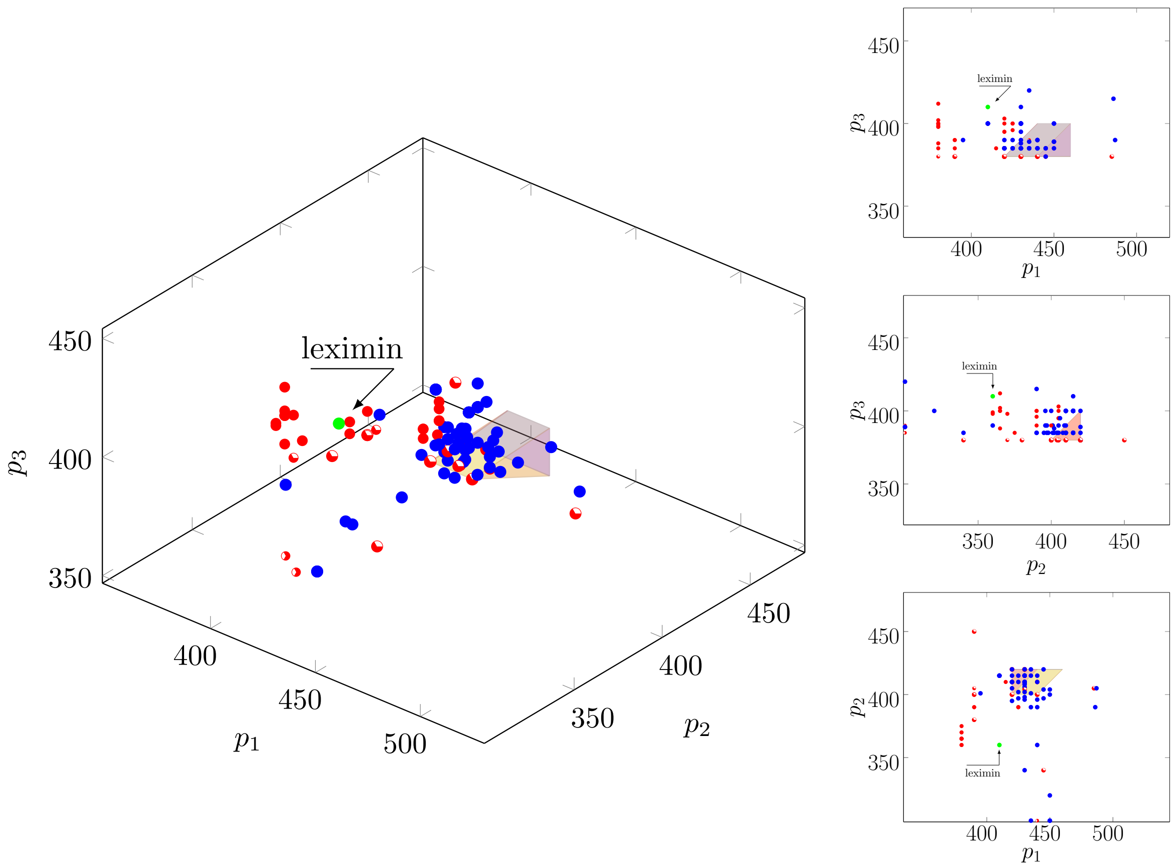

(c) Pit Trading

Figure 5: Resulting prices across treatments

Note: blue points are rounds with maximized total utility, filling of the markers corresponds to items traded, size of the markers is the total utility. When item is not traded, the minimum price is used instead of the trading price. Green point is the leximin price vector