1

The Exposure Problem and Market Design

JACOB K. GOEREE

AGORA Center for Market Design, UNSW, 2052 Sydney, Australia
E-mail: jacob.goeree@gmail.com

LUKE LINDSAY

University of Exeter Business School, University of Exeter, Exeter, EX4 4PU, UK
E-mail: luke.lindsay@gmail.com

Markets have an exposure problem when getting to the optimal allocation requires a sequence of transactions which if started but not completed leaves at least one trader with losses. We use laboratory experiments to evaluate the effect of the exposure problem on alternative market mechanisms. The continuous double auction performs poorly: efficiency is only 20% when exposure is high and 55% when it is low. A package market effectively eliminates the exposure problem: in low and high exposure treatments efficiency is 82% and 89% respectively. Building on stability notions from matching theory we introduce the concept of mechanism stability. A model of trade that combines mechanism stability with noisy best responses and imperfect foresight explains the difference in market performance. Finally, decentralized bargaining with contingent contracts performs well with perfect information and communication but not in the more realistic case when traders' preferences are privately known.

Keywords: *Exposure problem, package markets, market design, laboratory experiments*
JEL codes: *C78, C92, D47*

1. Introduction

Monetary theorists since at least Jevons have recognized how using money as a medium of exchange can facilitate trade. Ostensibly, the double coincidence of wants problem that occurs in barter can be solved by letting traders arrive at their desired allocation of goods via a series of bilateral transactions involving money. But when the market is thin and getting to a desired allocation requires a series of trades, the first of which leaves an agent worse off than not trading, the agent may be reluctant to make the first trade for two reasons. First, subsequent trades may not be executed. Second, even if it were certain that subsequent trades will occur, the initial trade may weaken the agent's bargaining position to the extent that the loss cannot be recouped. Either way, while the introduction of money solves Jevons' double coincidence of wants problem it does not protect traders from being exposed to losses. Anticipating this *exposure problem*, traders may be unwilling to make the first trade leaving potential gains from trade unrealized.

The goal of this paper is to examine how different market mechanisms perform in reassignment problems when exposure is present. The first mechanism we test is the continuous double auction (CDA). Our interest in the CDA is natural since it is the most-commonly used institution for contemporary financial and commodity markets. Furthermore, the CDA has an impressive track record in the lab and many experimenters would probably guess it would perform well in the simple environments we study: four subjects each own a house, each demands one house, and each has values for all four houses. When subjects' values are common knowledge, the possible gains from trade are apparent. Nevertheless, observed efficiency levels in the CDA are very low with many instances of no-trade or losses. While this poor performance contrasts with that of previous studies, it has an intuitive explanation in terms of exposure. In our setup, houses are substitutes, which implies that initial trades often result in losses. Traders risk being financially exposed when such losses cannot be recouped in subsequent trades, e.g. when there is strategic uncertainty about others' bargaining behavior.

To quantify the effects of exposure, we compare market performance in two parallel treatments. In the low-exposure treatment, all house values are shifted downward by a common constant compared to the values used in the high-exposure treatment. As a result, the optimal allocation and the total gains from trade are the same but the risk associated with buying a second house is less. We find that this manipulation has a strong positive effect. Efficiency levels are significantly and substantially higher in the low-exposure treatment, providing evidence for the impact of exposure on market

performance.[1]

The second mechanism we test is a package market that is a simple extension of the CDA. Like the CDA, it allows for standard buy and sell offers involving a single house and some amount of cash. In addition, it allows for arbitrary "package offers" involving several houses and cash, such as where one house is offered, one is demanded, plus some amount of cash is offered or demanded. Such package offers allow subjects to exchange houses without risking ending up with two houses or no house. And, unlike the top-trading-cycle procedure discussed below, such exchanges may involve money. The package market performs better than the CDA: efficiency is 82% when exposure is low and 89% when exposure is high.

The third mechanism we test is decentralized bargaining. We ran experiments to test whether the good performance of the package market can be achieved by decentralized trading. We find that decentralized bargaining with contingent contracts can deliver comparable efficiency levels to the package market when there is perfect information and communication is allowed. When house values are privately known, however, bargaining performs worse irrespective of whether communication is possible.

To put the experimental performance results in perspective, we simulate efficiency numbers for the well-known top-trading-cycle procedure (Shapley and Scarf, 1974).[2] Without money this simple procedure obviously cannot be fully efficient but it does outperform the CDA in both the low and high-exposure treatments. We also consider a variant of the ascending clock auction. Like the top-trading-cycle procedure, the modified ascending clock auction (MACA) is a strategy-proof mechanism that guarantees homeowners will end up at least as well off as their initial allocation.[3] The cost of this guarantee is that the mechanism does not always result in efficient allocations. In simulations, the MACA also outperforms the CDA.

Among the mechanisms tested, the package market performs best in the face of exposure: efficiency levels are high and significantly above those for the CDA. This improvement can partially be understood by comparing allocations that are stable under the two mechanisms. We say an allocation is m-stable if all allocations that can be reached via a single trade under mechanism m make at least one trader worse off. For example,

1. Another potential source of inefficiency is the fact that traders have complete information about who owns what house. In particular, they know when others are in a weak bargaining position, e.g. when holding two houses, which may create a hold-out problem. We find that revealing less information about who owns which house and previous trades reduces but does not eliminate efficiency losses.

2. The top-trading-cycle procedure proceeds in several steps: in each step, agents point to the house they prefer most among those available and houses (and owners) that form cycles are removed. A cycle may consist of a single owner pointing to their own house. A variant of the top-trading-cycle procedure is used for kidney exchange, see Roth, Sönmez, and Ünver (2005).

3. This mechanism was suggested to us by Philippe Jehiel.

in the CDA, an efficient swap of houses requires two trades and the status quo is stable if the first trade lowers total surplus. In contrast, in the package market, an efficient swap can be completed in a single trade so the status quo is not stable. More generally, assuming trade does not occur if the current allocation is stable predicts efficiency levels of 23% (70%) in the CDA when exposure is high (low). For the package market, predicted efficiency is 100% in both cases as an efficient reassignment is always possible via a single multilateral trade.

While m-stability produces aggregate efficiencies similar to observed levels, its deterministic predictions are trivially refuted by the individual trade data. Moreover, m-stability assumes myopic agents who think only one trade ahead. Building on recent approaches to "bounded rationality" we explore a more flexible model that can be estimated using individual trades. We consider agents who plan $k = 1, 2, \ldots$ steps ahead, akin to the level-k approach (e.g. Stahl and Wilson, 1994; Nagel, 1995), and who make noisy best responses, as in the QRE approach (e.g. Goeree, Holt, and Palfrey, 2016). Fitting this model to individual trade data reproduces the main features of the data including the improved efficiency of the package market relative to the CDA.

Since the package market is a straightforward adaptation of the CDA, it could potentially be applied in a variety of contexts. Besides real-estate, one could think of markets for other expensive durables such as cars, boats, etc. Another obvious candidate is financial markets where "pure swaps," i.e. package orders that do not involve money, are often introduced to mitigate the exposure problem. A different application concerns the trading of sports players. Whether a team wants to sell a certain player will often depend on whether they can find a suitable replacement. In these applications, package orders could facilitate more efficient outcomes especially when the market is thin.

1.1. Related literature

This paper contributes to an emerging literature on package markets, which builds on three more established strands: that on the continuous double auction, that on two-sided matching without money, and that on package auctions. Figure 1 shows the connections between the different mechanisms.

The continuous double auction: Vernon Smith's (1962) finding that behavior in the CDA robustly converges to competitive equilibrium outcomes is remarkable in that convergence occurs when it is not predicted. The experiments employ only a small number of buyers and sellers, there is no common knowledge of supply and demand, and subjects are not price takers but rather price makers. In these early experiments, however, exposure is

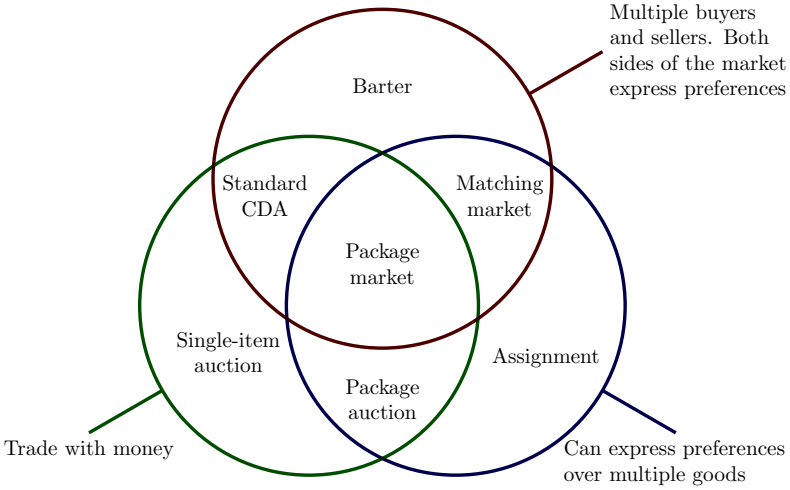

FIGURE 1
Relationships between mechanisms

not present. A few more recent studies have found limits to the domain where the CDA performs well. Van Boening and Wilcox (1996) find that the CDA fails in the presence of avoidable costs with observed efficiencies of 50% or less and highly erratic price dynamics. Mestelman and Welland (1987) find lower efficiencies with advance production compared to production on demand. One explanation for the CDA's poor performance in these settings is the effect of exposure. Our paper identifies a new simple setting where the CDA performs poorly and provides evidence that the poor performance is indeed due to exposure. The package market we propose restores efficiency by adding conditional offers to the CDA that protect traders from exposure.

Two-sided matching markets without money: The past two decades have seen important advances in the theory and application of matching mechanisms, e.g. assigning doctors to hospitals (Roth and Peranson, 1999) and matching kidney donors with recipients (Roth, Sönmez, and Ünver, 2004). Using mechanisms where participants can express preferences over multiple outcomes protects them from various forms of exposure. For example, with decentralized applications, newly trained doctors face exposure when hospitals offer placements with short deadlines. Should they accept an offer in hand and risk missing out on getting a better one later or let it expire and risk a worse outcome? A donor-recipient pair faces exposure when donating a kidney without simultaneously receiving

one in return.[4] Mechanisms based on Deferred Acceptance (Gale and Shapley, 1962) and Top Trading Cycles (Shapley and Scarf, 1974) provide elegant solutions to these problems when using money is not allowed. In settings where it is, however, they leave potential gains from trade unrealized. The package market we introduce takes one of the desirable features of matching mechanisms, i.e. allowing participants to express preferences over multiple goods to avoid exposure, and uses it in a mechanism with money so that the full gains from trade can be realized.

Package auctions: In one-sided auctions, the exposure problem arises when complementary goods are sold individually. A prominent example is the sale of spectrum licenses for wireless and mobile phone applications. Telecom operators typically want consecutive blocks of spectrum within a band or combinations of licenses that span adjacent geographic areas. In the simultaneous ascending auction, bidders compete for large numbers of individual licenses over a series of rounds, with provisional winners being announced after each round. This approach was pioneered by the US FCC in 1994 and has been copied in other countries with considerable success. But theoretical analyses (Goeree and Lien, 2014) and experimental evidence (e.g. Brunner, Goeree, Holt, and Ledyard (2010)) indicate that efficiency and revenue may be suppressed when bidders hesitate to incorporate synergistic values into their bids for fear they win only part of a desired combination. Package auctions avoid such exposure problems by allowing bidders to compete for combinations of items using "all-or-nothing" bids. The potential to improve efficiency and revenue has raised considerable interest in package auction design. Furthermore, several innovations proposed in the literature, e.g. the combinatorial clock auction, hierarchical package bidding, and sealed-bid combinatorial auctions, have been applied in recent spectrum sales (see Bichler and Goeree, 2017, for an up-to-date overview).

Package markets: There are several important differences that make the design of package markets much harder (Milgrom, 2007). Innovations in package auction design are unlikely to readily apply.[5] For example, in an auction setting, it is possible to design efficient, deficit-free mechanisms whereas in the market setting, it is generally not, see Loertscher, Marx, and Wilkening (2015) for a recent review. In the auction setting, it is possible to use a payment rule that, given reported values, selects prices from the core (Day and

4. To avoid this exposure problem, when several transplant operations are necessary they are conducted simultaneously.

5. See, however, Lubin, Juda, Cavallo, Lahaie, Shneidman, and Parkes (2008) who develop a package market built around the combinatorial ascending auction. The allocation and prices are determined iteratively with traders revising their orders at each step.

Milgrom, 2007);[6] in the market setting, the core does not exist for all reported values, so such a payment rule cannot be used. Finally, in a package auction, transactions are bilateral (between the auctioneer and one buyer), while in a package market, transactions can be multilateral (multiple buyers and/or multiple sellers).

Research on using package bidding in two-sided settings is much less developed. When multiple buyers and multiple sellers compete and both sides of the market value the items being traded, the exposure problem can arise with any type of good, not just with complements. The intuition is that even when goods are substitutes there can be complementarities between trades, as is the case for the house market studied here.

One approach is a direct mechanism or call market where participants submit orders once, and after a predetermined time, the allocation and prices are determined. Bossaerts, Fine, and Ledyard (2002) suggest a market of this form for trading securities when investors are interested in holding certain portfolios. Allowing traders to submit package orders protects against being left holding an unbalanced portfolio, which might otherwise occur when the markets are thin. Milgrom (2009) proposes a generalized message space – the space of assignment messages – for use in markets and other direct mechanisms where goods are substitutes. Our approach is different in that we extend a commonly-used market mechanism, the CDA, to accommodate package orders. This extension generalizes package auctions to the case with multiple buyers and multiple sellers with both sides of the market submitting preferences.

1.2. Organization

This paper is organized as follows. Section 2 presents definitions related to exposure and Section 3 describes the trading environment. In Section 4 we provide a detailed account of how the simple continuous double auction market, the package market, and decentralized bargaining are implemented. The experimental design is explained in Section 5. We next provide results on market efficiency (Section 6.1), the effect of exposure (Section 6.2), and then present the bargaining results (Section 6.3). In Section 7 we develop and estimate a Markov model of trading. Section 8 concludes. The appendix contains simulations with strategy-proof mechanisms (Appendix A), additional discussion of the bargaining results (Appendix B). Screenshots of the interface subjects used and sample instructions are in an online appendix.

6. Core pricing is used in the combinatorial clock auction, which has been used to sell spectrum in a number of countries since 2008.

2. The exposure problem

Consider an exchange economy with a set of agents I, a set of indivisible commodities H, and money. Agent $i \in I$ has quasi-linear utility $u_i(\boldsymbol{\omega_i}) + c_i$ where the pair $(c_i, \boldsymbol{\omega_i})$ is i's allocation with c_i the amount of money held and $\boldsymbol{\omega_i} \in \mathbb{Z}_{\geq 0}^{|H|}$ a vector of commodities. Agent i's initial allocation is denoted $(c_{i,0}, \boldsymbol{\omega_{i,0}})$. Agents can make a finite sequence of trades, labeled $t = 1, \ldots, T$, or not trade at all ($T = 0$). Trade t consists of the pairs $(y_{i,t}, \boldsymbol{x_{i,t}})$ for $i \in I$, describing the change in cash, $y_{i,t}$, and the change in commodities, $\boldsymbol{x_{i,t}}$, such that $\sum_{i \in I} y_{i,t} = 0$ and $\sum_{i \in I} \boldsymbol{x_{i,t}} = 0$. Agent i's allocation following trade t is $(c_{i,t}, \boldsymbol{\omega_{i,t}}) = (c_{i,0} + \sum_{j=1}^{t} y_{i,j}, \boldsymbol{\omega_{i,0}} + \sum_{j=1}^{t} \boldsymbol{x_{i,j}})$ and i's final allocation is $(c_{i,T}, \boldsymbol{\omega_{i,T}})$. The sequence of allocations can be used to define different aspects of exposure.

Definition 1. *An agent falls prey to exposure if their final allocation yields less utility than one of the previous allocations.*

Clearly, if agents can foresee the trading opportunities they will face, they should not fall prey to exposure. However, if an agent makes a series of trades and the prices of later trades are not fixed in advance, the agent may be exposed (at risk of falling prey to exposure).

Definition 2. *An agent makes an exposed trade if the allocation after the trade yields less utility than the allocation before the trade.*

Making an exposed trade does not imply falling prey to exposure. Indeed, getting to a competitive equilibrium allocation could involve an exposed trade.

A market mechanism m specifies the types of trades that are permissible. For example, whether trades involving multiple commodities or more than two parties are possible.

Definition 3. *There is an exposure problem in an economy with market mechanism m if there exists an allocation from which getting to the optimal allocation requires at least one trader to make an exposed trade.*

This definition allows us to determine whether a market mechanism has an exposure problem for a given economy. For example, consider a market mechanism where items are traded one at a time so the first trade involves agent i buying a single item from

agent j for some $i \neq j$. The gains π_i and π_j are defined as follows:

$$\pi_i = u_i(\boldsymbol{\omega_{i,1}}) - u_i(\boldsymbol{\omega_{i,0}}) - p$$

$$\pi_j = u_i(\boldsymbol{\omega_{j,1}}) - u_i(\boldsymbol{\omega_{j,0}}) + p$$

where p is the transaction price. If $\pi < 0$, then the agent makes an exposed trade. Clearly, if $\pi_i + \pi_j < 0$, then at least one agent makes an exposed trade. The quasi-linearity assumption implies that $\pi_i + \pi_j$ is independent of p, so transactions where one agent must make an exposed trade can be identified by only considering the item traded. Suppose the initial allocation is not optimal. Finding a sequence of non-exposed trades from the initial allocation to the optimal allocation establishes that there is not an exposure problem. One way to establish that there is an exposure problem is by showing all the first trades are exposed. Such allocations are stable in the following sense.

Definition 4. *An allocation is m-stable if no other allocation can be reached under mechanism m without at least one trader making an exposed trade.*

The next two sections introduce the economy and market mechanisms we study.

3. The reassignment game

In Shapley and Shubik's (1971) assignment game, there are m sellers and n buyers. Each seller is endowed with an item. The buyers value all items while the sellers value only the item they are endowed with. We study a symmetric variation of this game where all n agents play the role of both buyer and seller. Indivisible and differentiated items, houses, are traded for money. Each agent owns one house, so $|I| = |H|$. Agent i is initially endowed with house i.

Each agent demands exactly one house. Each agent has a private value for each of the houses, $v_i^h \sim U[\underline{v}, \bar{v}]$ where $0 \leq \underline{v} < \bar{v}$. Agent i's utility is $\max(v_i^1\omega_i^1, \ldots, v_i^n\omega_i^n) + c_i$. Let $\Omega = \{\omega_1, \ldots, \omega_n\}$ and Ω^* be the allocation of houses to agents that maximizes overall surplus. For this simple exchange economy, competitive prices always exist and are usually not unique. All the competitive prices support the efficient allocation and the set of competitive prices forms a bounded lattice (see also Shapley and Shubik, 1971).

An example with four agents is shown in Table 1. The numbers in the table represent agents' values for each of the houses. The underlined values indicate which house each agent is initially endowed with while the starred values indicate the allocation

TABLE 1

House values example

	Agent 1	Agent 2	Agent 3	Agent 4
House A	$\underline{60}^*$	34	59	36
House B	64	$\underline{31}$	57	43^*
House C	65	67^*	$\underline{68}$	43
House D	48	32	57^*	$\underline{34}$

Notes: Example of agents' values with four agents and four houses. The underlined values correspond to the initial allocation and the starred values to the optimal one.

that maximizes surplus. It is readily verified that the lower bound on the lattice of competitive prices is $(\underline{p}_A^* = 2, \underline{p}_B^* = 6, \underline{p}_C^* = 11, \underline{p}_D^* = 0)$ and the upper bound is $(\bar{p}_A^* = 39, \bar{p}_B^* = 43, \bar{p}_C^* = 67, \bar{p}_D^* = 37)$. Notice that although agent 3 starts with her most preferred house, trading to the optimal allocation at competitive prices does not make her worse off and can, depending on which vector of competitive prices is used, make her better off.

Despite the existence of a range of competitive equilibrium prices, the exposure problem may preclude efficient trade. Suppose houses are traded one at a time. To get to the optimal allocation, a series of trades is required. Consider the values shown in Table 1 and suppose the series starts with agent 2 buying house C from agent 3 at some price p_C. Agent 2's gain in utility is $\max(v_2^B, v_2^C) - v_2^B - p_C$ and agent 3's gain is $p_C - v_3^C$. The sum of the agents' gains is $\max(v_2^B, v_2^C) - v_2^B - v_3^C = \max(31, 67) - 31 - 68 = -32$. Since this sum is negative, whatever price the house was traded at, at least one of the agents must have made an exposed trade.

4. Trading mechanisms

This section describes the three trading mechanisms we evaluate: the simple CDA market, the package market, and decentralized bargaining. (The two strategy-proof mechanisms we consider are described in Appendix A.) In all the mechanisms, trade is voluntary. In both markets, traders submit orders in continuous time and trade occurs instantly when a set of compatible orders has accumulated. The markets differ in the types of order that are admissible. In the simple market, buy and sell orders are allowed; in the package market, buy, sell, and package orders are allowed. Under decentralized bargaining, traders propose contracts and a trade occurs when all the relevant parties accept a contract.

The following framework is used to describe traders' orders and holdings. An order

is a pair $o = (b, \boldsymbol{x})$ where b is a real number representing the amount of cash being offered or requested and $\boldsymbol{x} \in \{-1, 0, 1\}^N$ is a vector indicating which houses are offered or demanded. Positive values indicate an item is demanded and negative values indicate that it is offered. For example $(-20, \langle 0, 1, 0, 0\rangle)$ indicates "I am willing to pay up to 20 for house B" and $(30, \langle -1, 0, 0, 0\rangle)$ indicates "I am willing to accept 30 or more for house A." Orders are submitted in continuous time. An order is active until it transacts or is withdrawn. Let O^t denote active orders at time t and let O_i^t denote the active orders submitted by trader i. Elements of O^t are denoted $o_j = (b_j, \boldsymbol{x_j})$. Let $\boldsymbol{\omega_i} \in \{0, 1\}^N$ denote the houses held by trader i and c_i the amount of cash held by trader i.

In the simple market, two types of order are allowed: buying orders ($b < 0$ and exactly one component of $\boldsymbol{x}$ is 1 and the rest are zero) and selling orders ($b > 0$ and exactly one component of $\boldsymbol{x}$ is -1 and the rest are zero). In the package market, package orders are allowed in addition to buying and selling orders. A package order is an order that involves more than one house. The only restriction on package orders is that something must be given and something must be taken. Swaps involving cash, such as $(30, \langle -1, 0, 1, 0\rangle)$, are allowed. So are offers to buy, sell or exchange multiple houses, e.g. $(-50, \langle 0, 1, 1, 0\rangle)$, $(60, \langle -1, -1, 0, 0\rangle)$ or $(0, \langle -1, 0, 1, 1\rangle)$.

Each time a new order is submitted, an algorithm is run that determines if any transactions will occur. The winning orders (and hence the houses that get reallocated) are selected by maximizing the cash surplus. The cash surplus is calculated using the quantities traders specify in their orders. (Note that since the cash surplus depends on submitted orders rather than preferences, it need not correspond to the economic surplus.) Let $d_j = 1$ if order j is winning and $d_j = 0$ otherwise. The vector $\boldsymbol{d}$ is found by solving the following:

$$\max_{\boldsymbol{d}} \sum_{j \in O^t} -b_j d_j$$

subject to

indivisibility: $d_j \in \{0, 1\}$ for all $j \in O^t$

supply equals demand: $\sum_{j \in O^t} x_j^k d_j = 0$ for all $k \in H$

no short selling: $\omega_i^k + \sum_{j \in O_i^t} x_j^k d_j \geq 0$ for all $k \in H, i \in I$

budget constraints: $c_i + \sum_{j \in O_i^t} b_j d_j \geq 0$ for all $i \in I$

Let the set of winning orders be denoted $W = \{j \in O^t \mid d_j = 1\}$ and the set of losing orders $L = O^t \setminus W$. For losing orders, the submitter does not pay or receive anything. For winning orders, the submitter receives or pays an amount of cash $y_j \geq b_j$. In cases where $\sum_{j \in W} -b_j = 0$, the total amount of cash offered exactly matches the amount requested, so $y_j = b_j$. In cases where $\sum_{j \in W} -b_j > 0$, there is a cash surplus. No revenue is extracted, the entire cash surplus is redistributed. This means that for some $j \in W$, $y_j > b_j$. To determine the division of this cash surplus, a vector of prices $\boldsymbol{p}$ is chosen that solves the following:[7]

$$\boldsymbol{p} \cdot \boldsymbol{x_j} + b_j \leq 0 \text{ for all } j \in W$$

$$\boldsymbol{p} \cdot \boldsymbol{x_j} + b_j \geq 0 \text{ for all } j \in L$$

Once prices have been chosen, the payment for order j is $\boldsymbol{p} \cdot \boldsymbol{x_j}$.

An example of how the algorithm operates in the simple market is shown in the left panel of Table 2. The columns headings use the variables defined above. Each row in the table represents an order. Order 1 is offering to sell house A for 20. Order 2 offers to buy house A for 30 and order 3 offers to buy it for 27. The cash surplus is maximized if orders 1 and 2 are winning. A price for house A of 27 maximizes the minimum surplus subject to the constraint that supply equals demand.

The right panel of Table 2 shows an example for the package market. Order 1 offers to trade house B for house A without any money changing hands (a "swap"). Order 2 offers to trade C for house B and pay 6 in cash. Order 3 offers to buy house A and order 4 offers to sell house C. Finally, order 5 offers to swap house A for house C. There are two feasible sets of winning orders. First, a "three-cycle." consisting of orders 1, 2 and 5 which gives a cash surplus of 6. Second, a "chain" of length 3 consisting of orders 3, 4, and 5 which gives a cash surplus of 5. The three-cycle gives the higher cash surplus, so orders 1, 2, and 5 are winning and the cash surplus is divided evenly. Orders 1 and 5 receive 2 cash; order 2 pays 4 cash.[8]

In the two market institutions, traders submit orders. The orders are matched by an algorithm, which determines whether any transactions will occur and if so produces a contract that defines the terms of trade. One can think of a contract as a set of orders. In the bargaining institution, there is no centralized matching of orders. Instead, traders

7. Since the solution is not necessarily unique, a way to choose between alternatives is needed. The approach used is to lexicographically maximize the minimum surplus $y_j - b_j$, see Kwasnica, Ledyard, Porter, and DeMartini (2005).

8. When the winning orders involve more than one house, there is typically a range of house prices consistent with the cash payments. Hence, in contrast to the simple market, unique prices cannot be assigned to each house.

TABLE 2

Orders and transactions example

j	b	$\boldsymbol{x}$	d	y
1	20	$\langle -1,0,0,0\rangle$	1	27
2	-30	$\langle 1,0,0,0\rangle$	1	-27
3	-27	$\langle 1,0,0,0\rangle$	0	

j	b	$\boldsymbol{x}$	d	y
1	0	$\langle 1,-1,0,0\rangle$	1	2
2	-6	$\langle 0,1,-1,0\rangle$	1	-4
3	-25	$\langle 1,0,0,0\rangle$	0	
4	20	$\langle 0,0,-1,0\rangle$	0	
5	0	$\langle -1,0,1,0\rangle$	1	2

Notes: Examples of orders and transactions in the simple market (left) and package market (right).

propose contracts, and a trade occurs when all the relevant parties accept a contract. The only restriction on submitted contracts is that the budget must balance and no one gives anything they do not own.

The stability of allocations can be compared across the three mechanisms using the concept of m-stability. The package market and bargaining institution allow transactions between any two allocations, so non-optimal allocations are never package-market-stable or bargaining-stable. In contrast, the simple market only allows transactions where one house changes hands. Accordingly, there are non-optimal allocations that are simple-market-stable.

5. Experimental design

We conducted two sets of experiments to investigate the exposure problem in the 'reassignment game' described in Section 3. The first set compared the performance of the simple market and package market across a range of environments. A $2 \times 2 \times 2$ factorial design was used with the following factors.

Market design: The simple market was compared to the package market. This lets us test whether the exposure problem causes efficiency losses in the simple market and, if so, whether the package market performs better.

Level of exposure: A high exposure environment was compared to a low exposure environment. In the low exposure environment, house values were drawn uniformly from $[0, 50]$. In the high exposure setting, the draws were generated by adding 25 to the draws from the low exposure treatment. This increases the degree of exposure without changing the optimal allocation or the gains from trade. To see why exposure is worse, consider the sum of gains from the first trade where agent 2 buys house C from agent 3 (as in the example of Section 3). When 25 is added, the net gain is

$\max(v_2^B + 25, v_2^C + 25) - (v_2^B + 25) - (v_3^C + 25)$. Adding 25 to all the values reduces the gain from the first trade by 25. Accordingly, adding the constant tends to increase the number of exposed trades. Varying the degree of exposure lets us determine whether differences in market performance were caused by exposure or other factors.

Information structure: A complete information environment where subjects' values for the four houses were public information was compared to an incomplete information environment where subjects only knew their own values (and who owned which house). When values are public information, it is possible for agents to work out the optimal allocation and identify a sequence of trades to reach it. When values are private information, this is not possible. Accordingly, it is plausible that the exposure problem would cause greater efficiency losses under incomplete information. Varying the information structure lets us determine whether efficiency losses are caused by uncertainty about others' values or other factors such as strategic uncertainty and hold-out.

In the first set of experiments, the package market performed considerably better than the simple market. The second set of experiments aimed to answer some unresolved questions. In total, the second set of experiments included five new treatments.

Hiding exposed positions: A possible explanation for the poor performance of the simple market is hold-out. Subjects might be unwilling to take on two houses if others can see they have two houses as this weakens their bargaining position. To test this, an additional treatment with incomplete information and high exposure was run where who owned which house was hidden.

Bargaining and communication: Another natural question is whether the good performance of the package market could be replicated without the centralized processing of orders. To test this, four new treatments using decentralized bargaining were run. Treatments were run with both complete and incomplete information under high exposure. In these treatments, subjects proposed contracts involving two or more traders and specifying what each would give and take. If everyone involved in the contract accepted it, the contract was implemented immediately. Subjects could make as many proposals as they wished and could trade multiple times. In natural settings, bargaining usually involves negotiation, and in experiments, cheap talk often influences outcomes (see e.g. Crawford, 1998). It was not obvious what effect communication would have in our setting, so to give the bargaining institution the best chance of success, we ran treatments with and without communication. In treatments with communication, subjects could send freeform cheap-talk messages to other members of the group.

The following procedure was used in both sets of experiments. The instructions were read out loud to the subjects using a short PowerPoint presentation. During the presentation, subjects could ask questions in public. We chose this format to ensure common knowledge and to let us explain the user interface of the experimental software in detail.[9] After the instructions, there were three unpaid practice periods. This allowed subjects to gain experience of using the software and ask additional questions. The instructions and practice periods together typically lasted 30-40 minutes.

Subjects were assigned to groups of four people that were fixed for the rest of the experiment. There were 15 paid periods. In each period, subjects were endowed with a house and 100 cash. Subjects received new private value draws and endowments at the start of each period. Within a treatment, the draws varied across groups but the same draws were used across treatments (for example, trader 2 in group 1 in period 6 would have the same value draws in all treatments) to ensure the possible gains from trade were identical. In each period, there was three minutes of trading time.[10] In the market treatments, there was no limit on how many orders a subject could submit. Similarly, in the bargaining treatments, there was no limit on how many contracts a subject could propose. In the bargaining treatments with communication, periods lasted six minutes. During the first three minutes, the subjects could send messages to each other but not trade; during the remaining three minutes, they could send messages and trade.

A total of 312 subjects took part in the experiment (13 treatments with 24 subjects per treatment). There were two sessions for each treatment. Subjects were paid based on the realized gains from trade, i.e. for each subject in each period, earnings were calculated as $u(\text{final holdings}) - u(\text{endowment})$. The resulting values for each of the 15 periods were summed giving a total number of points earned in the experiment. Subjects were paid 0.2 Swiss Francs for each point plus a show-up fee. For the treatments without communication, the show-up fee was 15 Francs, average total earnings were 35 Swiss Francs and the sessions lasted 80 minutes. We used a higher show-up fee of 30 Francs for the treatments with communication because the longer periods meant the sessions took longer to complete. With communication, average total earnings were 55 Swiss Francs and the sessions lasted 120 minutes.

9. Screenshots of the software subjects used and the slides for the instructions are included in an online appendix.

10. In a pilot session, longer period times were tried. These produced similar results but subjects commented that the experiment was too slow.

6. Results

We compare the simple and package market institutions in terms of efficiency. We then discuss in detail how exposure affects the continuous double auction. Then we introduce and estimate a Markov model of trading. Finally, we consider whether decentralized bargaining with contingent contracts could solve the exposure problem.

6.1. Market performance

First, we focus on the proportion of the potential gains from trade that were realized in different treatments. Realized gains are calculated at the group level over the 15 periods:

$$\text{realized gains} = \frac{\sum_{t=1}^{15} U_t - \underline{U}_t}{\sum_{t=1}^{15} \overline{U}_t - \underline{U}_t} \times 100\%$$

where U_t is total surplus (the sum of the utilities of the four group members) in period t, $\underline{U}_t$ is the total surplus if there had been no trade, and $\overline{U}_t$ is the maximum possible total surplus. The gains realized in the different treatments are shown in Table 3. Consider the top panel of the table. Changes in the market mechanism or the degree of exposure have a clear effect on the proportion of gains realized, but whether or not subjects had complete information has no apparent effect. For this reason, the complete and incomplete information treatments are pooled in the rest of the analysis.

Result 1—Market design: In settings with exposure, more of the gains from trade are realized by the package market than the simple market.

In the high exposure setting, 20 percent of the gains from trade are realized in the simple market and 89 percent in the package market. Taking a group as the unit of observation, this difference is significant ($p < 0.001$, Mann-Whitney test, $n = 24$). In the low exposure setting, 55 percent of the gains from trade are realized in the simple market and 82 percent in the package market. Taking a group as the unit of observation, this difference is also significant ($p < 0.001$, Mann-Whitney test, $n = 24$). Similar patterns of results occur under complete and incomplete information. Three aspects of this result are remarkable. First, the low fraction of the gains from trade that are realized in the simple market. In other settings, the CDA often produces efficiency levels close to 100 percent. Second, the size of the effect of changing the market institution. In auction experiments, for example, different auction formats typically realize different proportions of the potential gains

TABLE 3

Realized gains from trade by treatment

	Incomplete information		Complete information		Pooled	
Exposure	Low	High	Low	High	Low	High
First set of experiments						
Simple market	57.4	19.7	53.2	19.7	55.3	19.7
	(6.8)	(8.8)	(6.0)	(10.3)	(4.4)	(6.4)
Package market	81.2	87.1	82.5	90.8	81.8	88.9
	(7.9)	(3.1)	(2.9)	(1.9)	(4.1)	(1.8)
Second set of experiments						
Hidden holdings		43.4				
		(6.0)				
Bargaining		60.6		78.8		69.7
		(9.3)		(4.3)		(5.8)
Bargaining + chat		62.2		90.6		76.4
		(12.7)		(2.4)		(7.7)
Simulations						
TTC	67.9 in all treatments					
MACA	71.6 in all treatments (61.0 excluding auctioneer)					

Notes: The percentage of the potential gains from trade that was realized in each of the 13 experimental treatments and the 2 simulations is shown. For the experimental treatments, bootstrap standard errors are shown in parentheses. These were calculated using 1000 bootstrap replications, taking a group as the unit of observation. The "Pooled" columns show averages of the "Complete information" and "Incomplete information" columns. The simulations are described in Appendix A. The simulations make the same predictions in all treatments because all treatments used the same value draws.

from trade. However, the differences are usually in the range of a few percentage points (e.g., Brunner et al., 2010). Third, the absence of a treatment effect when information about house values is made public. This indicates that observed inefficiencies are not due to information rents associated with private information but rather with strategic uncertainty about others' behavior.

A natural question is whether the package market only performs better in "difficult" cases where an exchange among three or four subjects is required to achieve the optimal allocation.

Result 2—Complexity: Market performance is not explained by the type of exchange cycle required to go from the initial to the optimal allocation.

We estimate the following linear model for each of the market types in each of the exposure settings

$$\text{realized gains}_{g,t} = \beta_1\, d[2]_{g,t} + \beta_2\, d[3]_{g,t} + \beta_3\, d[2,2]_{g,t} + \beta_4\, d[4]_{g,t} + \varepsilon_{g,t}$$

TABLE 4

Realized gains by complexity

	Simple low	Simple high	Package low	Package high
[2]	3.9	1.6	72.4	87.9
	(16.5)	(15.0)	(9.6)	(3.2)
[3]	44.7	12.4	81.1	88.9
	(6.7)	(9.2)	(4.7)	(2.3)
[2,2]	47.5	8.1	83.9	81.1
	(12.3)	(19.8)	(10.3)	(13.5)
[4]	39.6	1.0	74.2	75.1
	(13.8)	(23.6)	(4.8)	(7.2)
#clusters	12	12	12	12
n	172	172	172	172

Notes: There is one observation per group per period. Cases where the initial allocation was optimal are excluded. The dependent variable is the percentage of potential gains realized. The independent variables are dummies representing the complexity of the cycle that is needed to go from the initial to the optimal allocation. Standard errors are shown in parentheses and are adjusted for clustering at the group level.

The dependent variable is the percentage of potential gains realized. Each variable $d[C]_{g,t}$ is one for group g in period t if going from the initial to the optimal allocation involves cycle C (and it is zero otherwise). Here [2] indicates that going from the initial to the optimal allocation involves only a pair of subjects trading their houses. Similarly, $[2, 2]$ means that two such pairs are needed while [3] and [4] indicate cases where three or four subjects are needed to complete the exchange. The analysis is restricted to cases where the initial allocation is not optimal, hence exactly one of the $d[C]$ terms is one for each observation. There is no constant term. The estimates are shown in Table 4. For all four market-type and exposure combinations, the null hypothesis that $\beta_1 = \beta_2 = \beta_3 = \beta_4$ cannot be rejected ($p > 0.05$, F-test).

Result 2 shows it is not the complexity of the optimal trade cycle that drives the difference between the simple and package market. What does? There are two disadvantages to buying in the simple market. Since houses are substitutes the price paid for a second house typically exceeds the *increase* in value to the buyer, a loss that can be recouped only if the buyer is able to sell the first house. Second, owning two houses creates a weak bargaining position since the marginal value of the less preferred house is zero. Others may try to exploit this weaker position by waiting until the end of the period before making a low offer. Of course, foreseeing both types of problem, all group members may be hesitant to start trading and be the first to buy.[11] The next result suggests that the simple market is indeed prone to such "hold out" problems.

11. Note that these concerns do not apply when package orders are used since subjects can avoid owning two houses at any point in time.

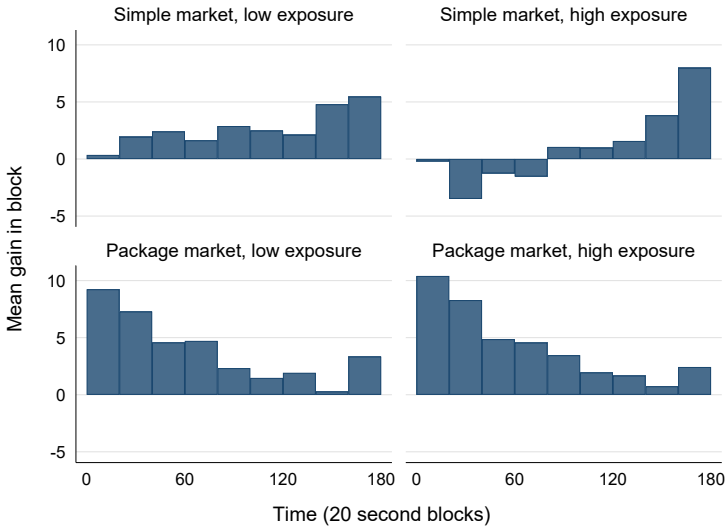

FIGURE 2
Evolution of realized gains from trade

Result 3—Holdout: In the simple market, most gains from trade are realized towards the end of the period. In contrast, in the package market, they are realized at the start of the period.

Figure 2 shows when gains or losses from trade occurred. The three-minute trading period is divided into nine 20 second blocks. The average number of points gained or lost during each block is shown for each of the treatments. Clearly, the simple CDA is subject to a severe holdout problem, which is virtually absent in the package market where most trading occurs in the first half of the period. Note from the top-right panel of Figure 2 that the simple market initially has negative gains from trade when exposure is high. In the next section, we investigate in more detail how exposure affects the performance of the CDA.

6.2. The effect of exposure

We now consider the effect of the level of exposure.

Result 4—Level of exposure: Decreasing the level of exposure raises the gains from trade in the simple market but not the package market.

In the simple market, 20 percent of the gains from trade are realized under high exposure and 55 percent under low exposure. Taking a group as the unit of observation, this difference is significant ($p = 0.002$, Mann-Whitney test, $n = 24$). Decreasing the level of exposure does not affect the gains from trade in the package market. Gains from trade fall from 89% to 82% but this difference is not significant ($p = 0.248$, Mann-Whitney test, taking a group as the unit of observation, $n = 24$). The difference between the high and low exposure treatments is that in the high exposure treatments all house values are 25 points higher. This means that the potential gains from trade are identical in both treatments but that losses from the first trade are larger in the high-exposure treatment.

The exposure problem can cause efficiency losses in two ways. Traders can fall prey to exposure by making exposed trades and not recouping losses. Alternatively, the prospect of falling prey to exposure can make traders reluctant to trade. The definition exposure (Section 2) can be used to identify cases where the exposure problem is present. If all the available first trades are exposed, then there is an exposure problem. The histograms in Figure 3 show the distribution of the gains and losses from the best first trade in the low and high exposure treatments. The figure shows how adding a constant to all values shifts the distribution of best first trades to the left. Notice that the shift does not change the shape of the distribution. The consequence of the shift is that there are fewer best first trades with a positive surplus, i.e. the exposure problem occurs more frequently.

Result 5—Exposed trades: When all the available first trades are exposed, the probability of no-trade and the probability of trade leading to losses both increase.

For the treatments that employed the simple market mechanism, when all available first trades involve a trader making an exposed trade, the frequency of no trade increases from 4.1% to 37.3% (5.7% to 40.3%). Similarly, the frequency of trade leading to losses increases from 6.1% to 28.8% (8.0% to 30.1%). These effects can be substantiated using Probit models:

$$\text{Prob}(\text{No trade} \,|\, x) = \Phi(\alpha + x\beta)$$

$$\text{Prob}(\text{Loss} \,|\, x) = \Phi(\alpha + x\beta)$$

There is one observation per group per round. If the best available first trade involves a loss, $x = 1$ and if not $x = 0$. Tables 5 and 6 show the results of estimating the two models with standard errors adjusted for clustering at the group level. When exposure is present, there is a significantly higher probability of no-trade and of the group making a loss. The losses typically resulted from failing to make additional trades after a loss-making first

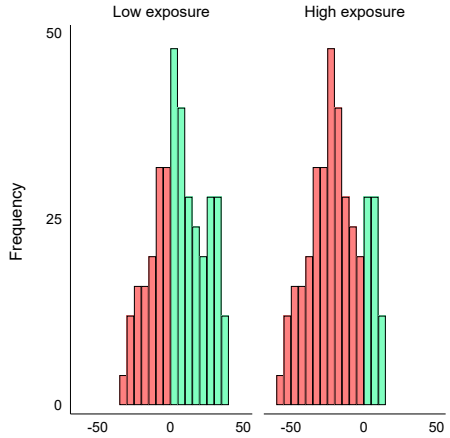

FIGURE 3

Histograms of the best first trades in the low and high exposure treatments

Notes: The dark bars correspond to negative best first trades, which indicate that at least one trader must make an exposed trade. When the best first trade gives a loss, the allocation is simple-market-stable.

TABLE 5

Probability of no trade

	No trade		
	Low	High	Pooled
Exposure	0.474^{***}	0.227^{**}	0.332^{***}
	(0.070)	(0.076)	(0.037)
# Groups	12	12	24
# Obs	180	180	360
Log likelihood	−59.62	−101.1	−165.1

Notes: Probit estimations of the probability of no-trade in the simple market using exposure as an explanatory variable. Marginal effects are reported. Standard errors of the marginal effects are shown in parentheses and are adjusted for clustering at the group level. * indicates $p < 0.05$, ** indicates $p < 0.01$, and *** indicates $p < 0.001$.

trade.

Figure 4 shows the initial and final unrealized gains from trade disaggregated by treatment. There is one point on the plot for each group in each period. Using the notation introduced earlier, the unrealized gains values were calculated as follows:

$$\text{Initial loss} = \underline{U}_t - \overline{U}_t$$

$$\text{Final loss} = U_t - \overline{U}_t$$

TABLE 6

Probability of losses

	Trade leading to loss		
	Low	High	Pooled
Exposure	0.248**	0.245**	0.227***
	(0.084)	(0.037)	(0.040)
# Groups	12	12	24
# Obs	180	180	360
Log likelihood	−70.25	−90.24	−161.1

Notes: Probit estimations of the probability of trade leading to losses in the simple market using exposure as an explanatory variable. Marginal effects are reported. Standard errors of the marginal effects are shown in parentheses and are adjusted for clustering at the group level. * indicates $p < 0.05$, ** indicates $p < 0.01$, and *** indicates $p < 0.001$.

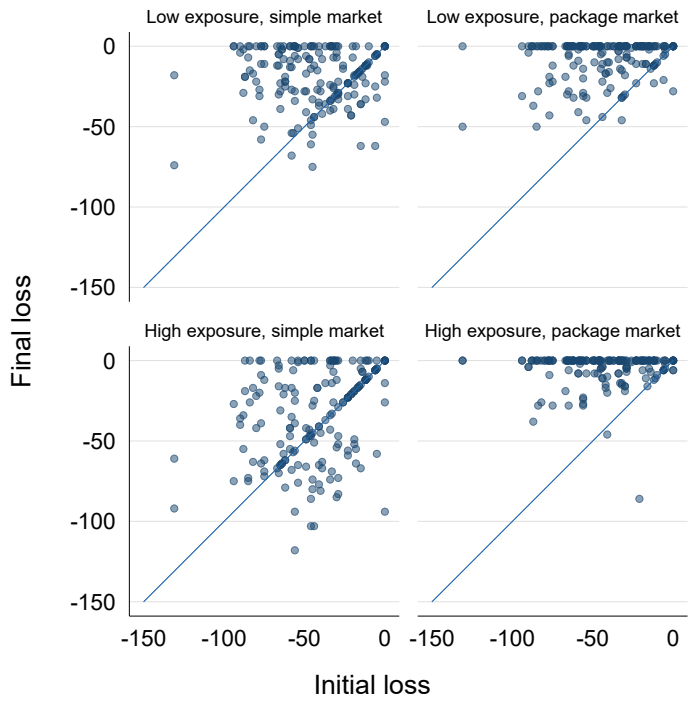

FIGURE 4

Realized and unrealized gains from trade in the simple market (left panels) and package market (right panels)

Notes: Points on the 45-degree line correspond to instances of no-trade and points below (above) the 45-degree line to instances of negative (positive) overall gains from trade.

This absolute measure of loss is used instead of a proportional one to make values from the high and low exposure treatments comparable. The vertical position of points on the

TABLE 7

Submitted orders

	Buy orders	Sell orders	Package orders
Simple low	35.2%	64.8%	–
Simple high	36.5%	63.5%	–
Package low	2.8%	11.7%	85.5%
Package high	2.2%	4.3%	93.6%

Notes: "Simple low" refers to the simple market with low exposure, "Package high" to the package market with high exposure etc. The three columns show the types of orders placed in the simple and package market under low/high exposure (with data from the complete and incomplete information treatments pooled).

graph indicates how much of the gains from trade were realized. A final loss of zero means all available gains from trade were realized. In all treatments, the optimal allocation was achieved by some groups in some periods. In periods where no trade occurred, points lie on the 45-degree line. This was common in the simple market and rare in the package market. Points below the 45-degree line indicate that there was trade but that the final allocation left the group worse off than they had started. Again, this occurred frequently in the simple market treatments and rarely in the package market.

The risk of exposure when buying did not go unnoticed by the subjects. The next result demonstrates that they mostly tried to sell their house rather than buying a second one and that those who bought were typically worse off.

Result 6—Submitted orders: In the simple market, sell orders are submitted more frequently than buy orders and those who sell first make significantly more than those who buy first. In the package market, package orders are submitted more frequently than simple orders.

Table 7 shows the percentage of buy, sell, and package orders disaggregated by treatment. In the simple market, it was not possible to submit package orders whereas in the package market, all types of order were admissible. In the simple market with high and low exposure approximately, two-thirds of the orders were offers to sell. This indicates that subjects were often unwilling to take on two houses. Indeed, subjects typically made more when they sold first (15.0 points and 13.0 points in the low and high exposure treatments respectively) than when they bought first (5.6 points and −4.8 points in the low and high exposure treatments respectively). The difference in gain between those who bought first and those who sold first is significant in the low and high exposure treatments ($p < 0.001$ and $p < 0.001$ respectively, Mann-Whitney tests). In the package

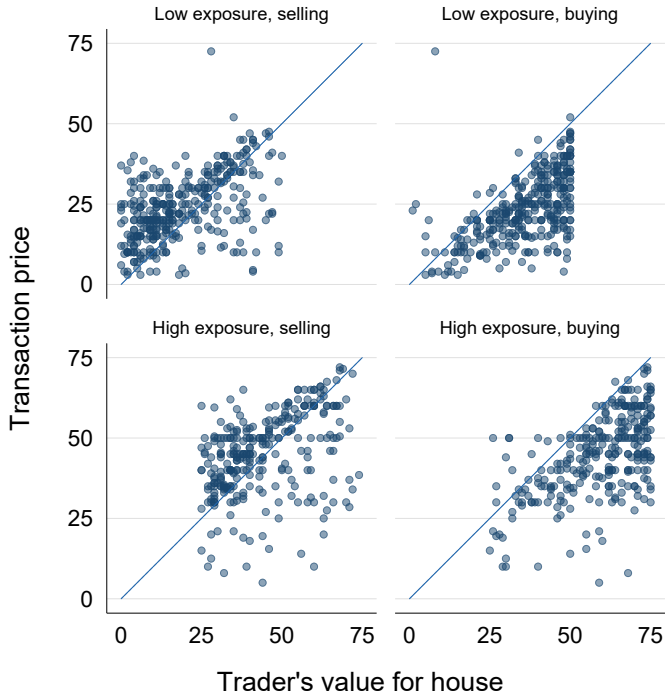

FIGURE 5
Scatter plot of transaction prices versus house values in the simple market with low exposure (top panels) and high exposure (bottom panels). The left panels show sell prices and the right panels show buy prices.

market, a large majority of subjects used package orders.

Figure 5 shows a scatter plot of transaction prices versus house values. The right panels indicate that subjects almost never paid more than their value for the house, which is to be expected if subjects act rationally. The sell prices shown in the left panels were frequently below value, which is not necessarily irrational. For example, when more than one house is held only the value of the best house counts, so selling one below value can be rational. Indeed, in 73 percent of the cases where the house was sold below value, the seller had two houses. In contrast, in only 28 percent of the cases where the house was sold above value did the seller have two houses. The difference is significant ($p < 0.001$, Pearson's chi-squared test). A natural question is whether the lower profits of traders who bought first was due to other traders being able to identify them and exploit their weak bargaining position. The first new treatment in the second set of experiments was

designed to disentangle the effect of this from other sources of inefficiency in the simple market.

Result 7—Hiding exposed positions: Hiding the holdings reduces but does not eliminate losses due to exposure.

In the simple market with high exposure and hidden holdings 43 percent of the gains from trade were realized compared to 20 percent when holdings were visible. This difference is significant ($p = 0.039$, Mann-Whitney test, taking a group as the unit of observation, $n = 18$). Table 3 shows the efficiency obtained in each of the treatments. The lower efficiency when holdings are visible is consistent with the conjecture that being seen holding two houses weakens one's bargaining position. When other traders cannot see you have two houses, you can sell for a higher price. However, the efficiency level of 43 percent achieved with hidden holdings is still substantially below the efficiency level of 89 percent achieved with the package market.

6.3. Bargaining

Two important features of the package market are the centralized matching of orders and the use of contracts where several houses change hands which protects traders against exposure. Could the good performance of the package market have been achieved by decentralized bargaining? The simple market imposes the constraint that houses are traded one at a time resulting in an exposure problem. Without this constraint, under complete information, one might expect bargaining to produce efficient outcomes. Four treatments in the second set of experiments explored this conjecture. Subjects traded using decentralized bargaining in the high exposure environment with complete and incomplete information and with and without freeform cheap-talk messages. The realized gains from the bargaining treatments are shown in the middle panel of Table 3.

Result 8—Bargaining and communication: Decentralized bargaining with contingent contracts only performs well under complete information. The effect of freeform communication is not discernible.

The difference between efficiency under complete and incomplete information is significant ($p = 0.011$, Mann-Whitney test, taking a group as the unit of observation, $n = 24$). In the bargaining treatments, allowing freeform communication seems to increase the realized gains but the effect is not statistically significant ($p = 0.184$, Mann-

Whitney test, taking a group as the unit of observation, $n = 24$). Although bargaining produces similar efficiency levels to the package market under complete information, it cannot replicate the performance of the package market in the more realistic setting with incomplete information.[12] This suggests that unless there is complete information and perhaps sufficient opportunity for communication, the centralized matching of orders provided by the package market is needed to achieve efficient allocations.

7. Markov trading model

This section develops a model of how the exposure problem affects market outcomes. We model the market as an absorbing Markov chain where states are allocations of houses to traders, moving between transient states represents trading and moving to an absorbing state represents trade ending. If agents never made exposed trades, an absorbing state would be entered upon reaching an m-stable allocation. Such a model, however, would be (trivially) refuted by the experimental results. Accordingly, we incorporate features of models with noisy best responses and strategic uncertainty.[13] This leads to less stark predictions and allows the parameters to be estimated from the experimental data. In the model, transition probabilities depend on how much traders gain from a trade. Two models of how traders think about the continuation game are considered. First, where agents only plan k trades ahead.[14] Second, where traders believe futures trades will only occur with probability q. A "precision parameter" λ determines how sensitive trades are with respect to gains. When $\lambda = \infty$ trade proceeds deterministically: until an m-stable allocation is reached if $k = 1$ or $q = 0$, but different degrees of foresight can be modeled by considering $k > 1$ or $q > 0$. In contrast, when $\lambda = 0$, behavior is random and all trades are equally likely. For intermediate values of λ, behavior is noisy, not deterministic. Agents do not always choose the best available trade although they do choose trades with higher gains more frequently. This allows the model to accommodate observed cases where agents fell prey to exposure. That is, where losses from an earlier trade were not fully recovered. The model is tractable and it allows us to make concrete predictions about the distribution of trades and final allocations.[15]

12. One reason is that agreements involving a single house are almost twice as prevalent in the incomplete information treatments. See Appendix B for details.

13. We thank the editor and an anonymous referee for suggestions that led us to develop this model.

14. The model has some similarities to level-k models. Beliefs are defined iteratively and higher values of k represent greater sophistication. In our model, k is the number of trades agents look ahead whereas in level-k models, it is the number of iterated best responses.

15. Modeling trading in the continuous double auction using standard game theory is challenging. There is a large action space, the move order is undefined, and actions occur in continuous time. This

The states in the Markov chain are modeled as follows. When there are n traders each endowed with one house, there are $n_\Omega = n^n$ ways to allocate the houses to traders. The allocations are denoted $\Omega_1, \ldots, \Omega_{n_\Omega}$ and the set of all allocations is denoted $\mathbf{\Omega_{all}}$. The Markov chain has a transient state and an absorbing state associated with each allocation. The reason for having two states associated with each allocation is to allow the number of trades to be endogenous and to allow trade to end at any allocation. The $2n_\Omega$ states are ordered such that all transient states appear before the absorbing states. Allocation Ω_r is associated with transient state X_r and absorbing state X_{r+n_Ω}. We can now define an adjacency matrix A. Entry $a_{rs} = 1$ if it is possible to transition from state X_r to state X_s and is zero otherwise. Transition is possible in the following cases. First, when the transition represents no trade. That is moving to an absorbing state ($s = r + n_\Omega$) or remaining in an absorbing state ($r = s$ and $r > n_\Omega$). Second, when the transition represents a permissible trade. Trades are transitions between transient states, that is when $r \neq s$, $r \leq n_\Omega$, $s \leq n_\Omega$. A trade is permissible if it is possible to get from the allocation Ω_r to allocation Ω_s. In the simple market, trades are only permissible if they involve a single house changing hands. In the package market, trades can involve any number of houses changing hands. Hence, the matrix A captures the differences between the simple market and the package market.

We assume that trades are more likely when agents believe they will yield a higher expected final surplus. Expected final surplus has two components. The immediate gain from the trade is described by π_{irs}. It denotes agent i's gain in utility from their holdings following the transition from X_r to X_s. The anticipated gains from the continuation game are described by matrix $\boldsymbol{\sigma}$. Entry σ_{is} represents agent i's belief about their gains in the continuation game after a transition to state X_s. Entries associated with absorbing states ($r > n_\Omega$) are zero. For transition X_r to X_s, the sum of the traders' gains is

$$\alpha^{\boldsymbol{\sigma}}_{rs} = \sum_{i \in I_{rs}} (\pi_{irs} + \sigma_{is})$$

where I_{rs} is the subset of agents whose holdings change. The probability of making a transition from X_r to X_s depends on the available transitions at r defined by the

makes a fully game-theoretic analysis almost certainly intractable. Without refinement such as sub-game perfection, many equilibria are possible in all market mechanisms. E.g. no one submitting orders is a Nash equilibrium; everyone submitting orders with competitive equilibrium prices is also a Nash equilibrium; Nash equilibria with less than full efficiency can be constructed by having some but not all traders submit orders. This would allow many outcomes to be rationalized but does not allow concrete predictions to be made. With sub-game perfection, perfect information, and an imposed predefined move order, backward induction should allow traders to execute sequences of trades that once completed leave all better off. This would produce efficient outcomes with or without package bidding.

adjacency matrix A, how α_{rs} compares to the value for other available transitions, and the precision parameter λ.

$$p_{rs}(\boldsymbol{\sigma}, \lambda) = \frac{a_{rs} e^{\lambda \alpha_{rs}^{\boldsymbol{\sigma}}}}{\sum_{t=1}^{2n_\Omega} a_{rt} e^{\lambda \alpha_{rt}^{\boldsymbol{\sigma}}}}$$

Note that transition probabilities are uniform when $\lambda = 0$ and deterministic when $\lambda \to \infty$. Using the function p_{rs} above, for a given $\boldsymbol{\sigma}$ and λ, a transition matrix $P_{\sigma\lambda}$ can be constructed. Entry p_{rs} is the probability of moving from X_r to X_s.

For simplicity, we assume that the price is chosen to split the gains from trade equally between traders, hence each trader receives $\frac{\alpha_{rs}^{\boldsymbol{\sigma}}}{|I_{rs}|}$. If beliefs are correct, then the following relation between the beliefs and transition matrix entries will hold for all agents i and for all states $r \leq n_\Omega$.

$$\sigma_{ir} = \sum_{s=1}^{2n_\Omega} p_{rs}(\boldsymbol{\sigma}, \lambda) \frac{\alpha_{rs}^{\boldsymbol{\sigma}}}{|I_{rs}|}$$

We do not impose the assumption that beliefs are correct. Instead, we allow the agents to only plan k trades ahead. Let $\boldsymbol{\sigma^1}$ denote the beliefs for $k = 1$. In this case, there are no further trades, so for all states $\sigma_{is} = 0$. Beliefs for $k > 1$ are then defined iteratively as follows.

$$\sigma_{ir}^{k+1} = \sum_{s=1}^{2n_\Omega} p_{rs}(\boldsymbol{\sigma^k}, \lambda) \frac{\alpha_{rs}^{\boldsymbol{\sigma^k}}}{|I_{rs}|}$$

We also consider a model were agents believe trade will continue with probability q after each trade.

$$\sigma_{ir} = q \sum_{s=1}^{2n_\Omega} p_{rs}(\boldsymbol{\sigma}, \lambda) \frac{\alpha_{rs}^{\boldsymbol{\sigma}}}{|I_{rs}|}$$

Using the equations above, for a given pair of parameters λ and k (or λ and q), a transition matrix can be produced. This matrix gives for each possible allocation, the probability of different trades occurring and the probability of trade ending. In addition, given a transition matrix P, there are established procedures for deriving a matrix of absorption probabilities B such that entry b_{rs} is the probability of eventually being absorbed into state s given the current state r.[16]

Table 8 shows the parameters, log-likelihood scores, and predictions of different versions of the Markov trading model. The transition matrix and the observed trades are used to calculate a log-likelihood score for the model. The transition matrix is also

16. The steps required to derive B from P are described in Grinstead and Snell (1997) chapter 11.

TABLE 8

Markov models

	Models with fixed parameters			Models with estimated parameters					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sample	Pooled	Pooled	Pooled	Simple	Package	Pooled	Simple	Package	Pooled
λ	0	∞	∞	0.061	0.115	0.089			
	0.063	0.115	0.089						
k	1	1	2	1	1	1			
q							0.096	0	0
Loglikelihood	-7653	$-\infty$	$-\infty$	-1775	-2067	-4018	-1773	-2067	-4018
				Predicted percentage of gains realized					
Simple, low	−54	70	100	50		63	52		63
Simple, high	−121	23	100	36		33	40		33
Package, low	−61	100	100		92	87		92	87
Package, high	−134	100	100		94	91		94	91
RMSD	160	13	46	12	8	7	14	8	7

used to predict the efficiency in the high and low exposure settings with the simple market and package market (the lower panel of the table). Finally, the root-mean-square deviation between the predicted efficiencies and observed efficiencies is calculated (the row 'RMSD').

For models 1-3, the parameters were chosen to explore their effect on the model's fit and predictions. In model 1, $\lambda = 0$ which means transition probabilities are independent of payoffs and so all trades are equally likely. In contrast, in models 2 and 3, $\lambda = \infty$ which means that the selected transition (either a trade or trade ending) is the one that, given beliefs, gives the highest payoff. In model 2, beliefs are based on looking one trade ahead which means that trades that lead to allocations with a lower value than the current allocation are never selected. The model predicts trade continues until a stable allocation is reached. An allocation is m-stable (see Definition 4) if there is no single trade that increases the sum of the traders' surplus. In the package market, all allocations can be reached in one trade so there is no efficiency loss. In the simple market, there are some stable allocations which are not efficient since getting to a more efficient allocation requires more than one trade. The pattern of predicted efficiencies is similar to what was observed in the experiment, but because the model is deterministic, it cannot account for the noise in the experimental data. In model 3, where $k = 2$ beliefs are based on looking two trades ahead, so temporary surplus losses are tolerated if the loss is recouped in the subsequent trade. This model predicts full efficiency in all treatments.

Maximum likelihood estimation of the Markov model's parameters leads to the following result.

Result 9—Markov model: For both the simple and package market, the best fitting Markov model is one where beliefs are based on planning one trade ahead.

For models 4-6, the parameters λ and k are estimated by maximum likelihood estimation. Model 4 is estimated using data from the simple market, model 5 using data from the package market, and model 6 pooling data from both mechanisms. For all three models, the estimated value of k is one. The predicted efficiencies are relatively close to the levels observed in the experiment. Similarly, for models 7-9, the estimated values of q are close to zero. This is consistent with agents planning one trade ahead.

One explanation for this finding is agents are bounded rational and do not plan more than one trade ahead. An alternative explanation is strategic uncertainty. The probabilities of other's actions in the continuation game are unknown. An agent who is ambiguity averse or reasons based on the worst-case scenario may only consider the immediate payoffs of a trade, which are certain, and ignore potential but uncertain gains in the continuation game.

8. Concluding remarks

The experiments reported in the paper were deliberately designed to be simple. Items were substitutes, there were well-defined property rights and no transaction costs. In addition, in half the treatments there was perfect information. These are conditions where one might expect the Coase theorem to hold and an efficient outcome to occur no matter how property rights are allocated.[17] The results show that in a standard double auction market only a small fraction of the total gains from trade are realized, both with complete and incomplete information. This poor performance is due to the exposure that arises when going from the initial allocation to the optimal one requires someone to temporally make a loss. The package market introduced in this paper largely solved the problem. By allowing for orders that include both a sell and a buy plus some amount of cash, the package market eliminates the exposure problem and produces efficient outcomes in situations where the continuous double auction and the top-trading-cycles procedure fail.

The package market shares some features with contingent contracting, which can also be used to reduce exposure.[18] For example, Collins and Isaac (2012) find that the

17. What has become known as the Coase Theorem was not presented as a theorem by Coase himself and the concept is somewhat nebulous. Parisi (2008) provides a modern interpretation: 'The Coase Theorem predicts that, in a competitive market environment without legal or factual impediments to exchange, the final allocation of rights will be efficient.' On this reading, one could argue that in the simple market the restriction that houses are traded one at a time is an impediment to exchange, and accordingly, the poor performance of the simple market is not contrary to the theorem.

18. Contingent contracts are used in a range of settings and can take various forms. Payments can be contingent on a natural event occurring, for instance flood insurance, or payments can be contingent on prices, for instance employment contracts with a wage indexed on the rate of inflation.

holdout problem in land assembly can be mitigated using contingent contracts. In some countries, real estate sale contracts can be contingent on the buyer selling their home, which removes the risk of being left with two houses. There are important differences with the proposed package market, however. First, in the context of the real-estate example, contingent contracts typically restrict the seller from selling to another buyer, in a sense shifting the exposure from the buyer to the seller, a feature that is not present in the package market. Second, the package market provides a flexible solution in that orders in the package market do not have to identify a counter party, e.g. an offer to exchange house A for house B does not specify who will take house A. The offer could be part of a transaction cycle of length three or more, in which case it is not the owner of house B that takes house A. Importantly, when submitting orders, traders do not have to worry about which type of transaction cycle will result. In our experiments, decentralized bargaining with contingent contracts delivers efficiency levels comparable to those of the package market if there is perfect information and communication is allowed. But in the more realistic case when house values are privately known, the package market outperforms contingent contracting.

The package format introduced in this paper is a simple extension of the continuous double auction. As such it has the promise to be applicable in a variety of circumstances where agents desire to complete all or none of a sequence of trades but there is uncertainty about whether some of them can be completed. Examples include markets for expensive durables, corporate bond markets, trading of sports players, and emission permits (Fine, Goeree, Ishikida, and Ledyard, 2017). Another example is the reallocation of airport resources. Landing and take-off slots are complements, so airlines would benefit from being able to bid for packages of compatible slots. In the long term, an airline may intend to expand the number of flights per day or number of destinations served. In the short term, adverse weather conditions such as thunderstorms can decrease an airport's capacity requiring slots to be reallocated (see Balakrishnan, 2007). A package market, with appropriate safety constraints, could help ensure slots get allocated efficiently. In less developed countries where small fragmented farm plots are common, a package market could help consolidate land holdings. Bryan, de Quidt, Wilkening, and Yadav (2017) conduct a framed field experiment among Kenyan farmers to compare the performance of several mechanisms for land exchange. They find that a variant of the package market we propose is able to achieve high levels of efficiency, yielding important gains to farmers.

A final example concerns the reallocation of licenses to use radio spectrum. Such licenses have been auctioned off by the US government since 1994. Over time, demand for services that rely on radio spectrum has changed and the technology to exploit

spectrum has improved, e.g. digital television requires much less bandwidth than analogue transmission. Furthermore, telecom operators that successfully participated in different spectrum auctions now typically own licenses that are dispersed both in the geographic and frequency domains. Since geographically adjacent, contiguous blocks of spectrum are more valuable there are likely gains from trade. A package market could facilitate a more efficient allocation of licenses while ensuring telecom operators that their overall network capacity remains intact.

APPENDIX

Appendix A. Simulations with strategy-proof mechanisms

We considered two strategy-proof mechanisms. First, the top-trading-cycles procedure – described by Shapley and Scarf (1974) but attributed to David Gale – that reallocates houses without cash transfers. Each house owner ranks the houses from best to worst. House owners point at the house they rank highest among those available (pointing at one's own house is allowed). When cycles form, the owners are assigned the house they are pointing at and the house and owner are removed. A house and owner is part of a cycle if following the path defined by the pointing leads back to the owners' house. The process is repeated with the remaining houses and owners until all have been removed.

Second, a modified ascending clock auction (MACA). This mechanism was suggested to us by Philippe Jehiel, who also provided his notes, joint with Olivier Compte, on the mechanism (personal communication, July 23, 2012). In a setting where initially houses are not allocated, it is possible to allocate them efficiently by running an ascending clock auction as described by Demange, Gale, and Sotomayor (1986). If houses are already allocated, running the standard ascending clock auction can make some participants worse off than if they kept their initial allocation. The modified ascending clock auction guarantees that participants will end up at least as well off as with their initial allocation. The cost of this guarantee is that the mechanism will not always give an efficient allocation.

In this mechanism, each agent is assigned one house, so allocations can be described by a mapping $\mu : I \mapsto H$. Let the initial assignment of houses to agents be given by μ_0. Let the initial vector of prices be $\boldsymbol{p_0}$ with $p_0(h) = 0$ for all houses. Let v_i^h be i's valuation for h. The pair (μ, p) specifies the house $\mu(i)$ that i gets and the price $p(h)$ paid for h. The participation constraint is i should get no less than $v_i^{\mu_0(i)}$.

The mechanism works as follows. In round t, the vector of prices is $\boldsymbol{p_t}$ and person i's demand is $D_i(p_t) = \arg\max_h(v_i^h - p_i^t(h))$ where

$$p_i^t(h) = \begin{cases} p_t(h) & \text{if } h \neq \mu_0(i) \\ 0 & \text{if } h = \mu_0(i) \end{cases}$$

If $\mu_0(i) \in D_i(p_t)$ for some i, then i gets $\mu_0(i)$ at price zero, house $\mu_0(i)$ and individual i are withdrawn and the process continues. Otherwise, if there are some over-demanded houses, their price is increased. Otherwise, the process stops, i gets $h \in D_i(p^t)$ and pays $p_t(h)$.

For each of the groups and each of the periods, the allocation that would be produced by running the Top-Trading-Cycles and Modified Ascending Clock Auction were found. The proportion of realized gains from running the TTC is 68 percent. For the MACA it is 72 percent although a proportion of this is revenue collected by the auctioneer. If the auctioneer‘s revenue is not included, the figure is 61 percent. For the simulations, it was assumed everyone plays their dominant strategy. Despite this, the efficient outcome is not always obtained. This is because obtaining the efficient allocation through voluntary trade sometimes involves one or more agents receiving monetary compensation for moving to a less preferred house. In the TTC and MACA mechanisms agents never end up in a less preferred house so the mechanisms cannot always achieve efficient outcomes. The efficiency figures are considerably less than the proportion of gains actually realized in the package market but considerably more than was

TABLE 9

Bargaining proposals and agreements

	Proposals % with # houses				Agreements % with # houses			
Treatment	1	2	3	4	1	2	3	4
Incomplete info.	32	55	4	9	23	77	0	0
Incomplete info. + chat	29	66	3	2	12	86	0	2
Complete info.	23	68	4	5	11	86	3	1
Complete info. + chat	15	71	9	4	6	83	7	4
All	28	61	4	7	12	83	3	2

Notes: In the bargaining treatments subjects submitted proposals specifying who would buy which house and the price. If all traders named in the proposal accepted, the proposal became an agreement and was executed. Proposals and agreements could involve 1-4 houses. The table reports the percentage of proposals/agreements involving the specified number of houses.

realized in the simple market.

Result 10—strategy-proof mechanisms: Top-trading-cycles and the modified ascending clock auction realize more of the gains from trade than the simple market but less than the package market.

The gains from trade realized in the package market are significantly higher than those that the two strategy-proof procedures could have achieved. A t-test rejects the null hypothesis that the realized gains from trade in the package market are equal to 68%, for both the low ($p = 0.009$) and high ($p < 0.001$) exposure treatments.[19] In contrast, the gains realized in the simple market are significantly lower than those that the two strategy-proof procedures could have achieved. A t-test rejects the null hypothesis that the realized gains from trade in the simple market are equal to 68%, for both the low ($p = 0.02$) and high ($p < 0.001$) exposure treatments. It is interesting that the simple top-trading-cycles procedure outperforms the CDA in both the low and high-exposure treatments. It should be noted, however, if the mechanisms had been run with human subjects, there may have been additional efficiency losses due to subjects not playing their dominant strategies. For instance, Chen and Sönmez (2002, 2006) find that in experiments, a significant proportion of subjects do not play their dominant strategies in the TTC mechanism.

Appendix B. Further details about the bargaining experiments

In this appendix we analyze the types of agreements that occur in the bargaining treatments, e.g. whether they involve multiple houses.

Table 9 shows the distribution of proposals and agreements involving different numbers of houses. Consider the columns showing the percentage of proposals and agreements involving one house in different treatments. Allowing communication and switching from incomplete to complete information appears to be associated with less use of one house agreements.

Result 11—Bargaining agreements: Agreements involving only a single house were twice as prevalent in the treatments with incomplete information, whether or not freeform communication was allowed. Likewise, agreements involving only a single house were twice as prevalent in the treatments without freeform communication, irrespective of whether values were private or commonly known.

19. When comparing simulation results to experimental results, there is only one random sample since the simulations are deterministic. Accordingly, we use a one-sample t-test instead of the two-sample Mann-Whitney.

Acknowledgements

We would like to thank the Swiss National Science Foundation (SNSF 138162), the European Research Council (ERC Advanced Investigator Grant, ESEI-249433), and the Australian Research Council (ARC-DP150104491 and ARC-DP190103888) for financial support. We are grateful for the useful suggestions we received from the Editor and referees as well as numerous colleagues and seminar participants.

References

Balakrishnan, H. (2007). Techniques for reallocating airport resources during adverse weather. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pp. 2949–2956. IEEE.

Bichler, M. and J. K. Goeree (2017). *Handbook of spectrum auction design*. Cambridge, UK: Cambridge University Press.

van Boening, M. V. and N. T. Wilcox (1996). Avoidable cost: Ride a double auction roller coaster. *The American Economic Review 86*(3), 461–477.

Bossaerts, P. L., L. Fine, and J. O. Ledyard (2002). Inducing liquidity in thin financial markets through combined-value trading mechanisms. *European Economic Review 46*(9), 1671–1695.

Brunner, C., J. K. Goeree, C. A. Holt, and J. O. Ledyard (2010). An experimental test of flexible combinatorial spectrum auction formats. *American Economic Journal: Microeconomics 2*(1), 39–57.

Bryan, G., J. de Quidt, T. Wilkening, and N. Yadav (2017). Land Trade and Development: A Market Design Approach. CESifo Working Paper Series 6557, CESifo Group Munich.

Chen, Y. and T. Sönmez (2002). Improving efficiency of on-campus housing: An experimental study. *The American Economic Review 92*(5), 1669–1686.

Chen, Y. and T. Sönmez (2006). School choice: an experimental study. *Journal of Economic Theory 127*(1), 202–231.

Collins, S. M. and R. M. Isaac (2012). Holdout: Existence, information, and contingent contracting. *Journal of Law and Economics 55*(4), 793–814.

Crawford, V. (1998). A survey of experiments on communication via cheap talk. *Journal of Economic Theory 78*(2), 286–298.

Day, W. R. and P. Milgrom (2007). Core-selecting package auctions. *International Journal of Game Theory 36*(3), 393–407.

Demange, G., D. Gale, and M. Sotomayor (1986). Multi-item auctions. *Journal of Political Economy 94*(4), 863–872.

Fine, L., J. K. Goeree, T. Ishikida, and J. O. Ledyard (2017). Ace: A combinatorial market mechanism. In M. Bichler and J. K. Goeree (Eds.), *Handbook of Spectrum Auction Design*. Cambridge, UK: Cambridge University Press.

Gale, D. and L. S. Shapley (1962). College admissions and the stability of marriage. *The American Mathematical Monthly 69*(1), 9–15.

Goeree, J. K., C. A. Holt, and T. R. Palfrey (2016). *Quantal Response Equilibrium: a stochastic theory of games*. Princeton University Press.

Goeree, J. K. and Y. Lien (2014). An equilibrium analysis of the simultaneous ascending auction. *Journal of Economic Theory 153*, 506–533.

Grinstead, C. M. and J. L. Snell (1997). *Introduction to probability* (Second ed.). American Mathematical Soc.

Kwasnica, A. M., J. O. Ledyard, D. Porter, and C. DeMartini (2005). A new and improved design for multiobject iterative auctions. *Management Science 51*(3), 419–434.

Loertscher, S., L. M. Marx, and T. Wilkening (2015). A long way coming: Designing centralized markets with privately informed buyers and sellers. *Journal of Economic Literature 53*(4), 857–897.

Lubin, B., A. Juda, R. Cavallo, S. Lahaie, J. Shneidman, and D. C. Parkes (2008). ICE: An Expressive Iterative Combinatorial Exchange. *Journal of Artificial Intelligence Research 33*, 33–77.

Mestelman, S. and D. Welland (1987). Advance production in oral double auction markets. *Economics Letters 23*(1), 43–48.

Milgrom, P. (2007). Package auctions and exchanges. *Econometrica 75*(4), 935–965.

Milgrom, P. (2009). Assignment messages and exchanges. *American Economic Journal: Microeconomics 1*(2), 95–113.

Nagel, R. (1995). Unraveling in guessing games: an experimental study. *American Economic Review 85*(5), 1313–1326.

Parisi, F. (2008). Coase theorem. In S. N. Durlauf and L. E. Blume (Eds.), *The New Palgrave Dictionary of Economics*. Basingstoke: Palgrave Macmillan.

Roth, A. E. and E. Peranson (1999, September). The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review 89*(4), 748–780.

Roth, A. E., T. Sönmez, and M. U. Ünver (2004). Kidney exchange. *The Quarterly Journal of Economics 119*(2), 457–488.

Roth, A. E., T. Sönmez, and M. U. Ünver (2005). A kidney exchange clearinghouse in new england. *The American Economic Review 95*(2), 376–380.

Shapley, L. S. and H. E. Scarf (1974). On cores and indivisibility. *Journal of Mathematical Economics 1*(1), 23–37.

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

Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy 70*(2), 111–137.

Stahl, D. and P. Wilson (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization 25*(3), 309–327.