Avoidable Costs and Market Design

Jacob K. Goeree, Luke Lindsay, and Xavier Del Pozo*

August 27, 2026

Abstract

Virtually all producers, most notably airlines and electricity generators, face avoidable fixed costs that cause non-convexities in production. Standard bilateral trading institutions such as the continuous double auction produce inefficient and unstable outcomes in the presence of avoidable costs. We design a package market that allows traders to submit schedules of quantity-contingent bids. In an experiment, we compare two continuous double-auction formats to two schedule-based variants. The new mechanisms substantially outperform the continuous double auctions, both in efficiency and in preventing seller losses. Efficiency is highest when schedules can be revised in continuous time. The mechanisms that allow sellers to express non-convex costs restore efficiency in markets where bilateral trading institutions fail.

Keywords: Market design, experiments, package markets, avoidable costs
JEL Classification: D47, L19

*Goeree: AGORA Center for Market Design and the School of Economics, UNSW Australia Business School, Sydney NSW 2052, Australia; Email: jacob.goeree@gmail.com. Lindsay: University of Exeter. Email: luke.lindsay@gmail.com. Del Pozo: Innovation Process Technology AG, Zurich. Email: xavier.delpozo@protonmail.ch. We are grateful for helpful comments from Brett Williams and from seminar audiences at Chapman University, UC Santa Barbara, UC Santa Cruz, and the University of East Anglia, and from conference participants at the 14th Conference on Economic Design (University of Essex), the ESA International Meeting (Hawaii), the ESA North American Meeting (Tucson), and the Frontiers of Bargaining: Theory and Experiments workshop (NYU Abu Dhabi). Goeree gratefully acknowledges funding from the Australian Research Council (DP220102893).

1. Introduction

In many industries, sellers must incur an avoidable cost before they can supply any units. Power plants and production lines have start-up costs; aircraft and ships must be dispatched, incurring fuel and crew costs; and entertainment venues must be prepared and staffed before opening. In each case, the cost can be avoided when output is zero, but must be incurred when output is positive. Once the cost is incurred, additional units can be supplied at low marginal cost.

Avoidable costs typically lead to declining average costs. In markets with heterogeneous capacities, this creates an exposure problem. A seller may be willing to supply a block of units only if sufficient demand materializes to cover the avoidable cost, but cannot ensure this outcome through sequential bilateral trades. Selling the first unit exposes the seller to losses if insufficient units are sold, while staying out forgoes efficient gains from trade. As a result, markets can exhibit low efficiency even when mutually beneficial trades exist. The underlying difficulty is that efficient allocations are inherently multilateral, while the trading process is bilateral.

The continuous double auction (CDA) has proven to be the most efficient and competitive trading institution in laboratory experiments. However, Van Boening and Wilcox (1996) showed that efficiency is seriously undermined in the presence of avoidable costs. Other approaches in these types of markets have not restored high levels of efficiency, including two-part price schemes (a fixed fee for production and a price per unit) and nonlinear pricing through bilateral contracting.

We design a package market to allow more efficient outcomes to be obtained in settings with avoidable costs. Several challenges prevented a standard market mechanism from being used. The non-convexity caused by avoidable costs means that competitive prices may not exist (Scarf, 1994), achieving an efficient outcome may require price discrimination, and the core may be empty (Telser, 1994). In our mechanism, buyers and sellers submit schedules of quantity-contingent bids, and then an allocation that maximizes surplus given the submitted bids is implemented. Individualized prices are chosen to be as close as possible to competitive prices, while ensuring that the mechanism is strictly budget-balanced and individually rational: no buyer pays more than their bid and no seller receives less than their ask.

We use laboratory experiments to assess how the new mechanisms perform relative to the continuous double auction. We compare two versions of the continuous double auction – CDA Hidden and CDA Visible in which the order book is either hidden or visible – to two versions of the package market – One-shot Schedules and Revisable Schedules in which schedules can be submitted only once or revised subject to an improvement rule. We use an environment where costs and values are redrawn each period rather than stationary repetition.

This paper makes three contributions to the literature on market design and avoidable costs. First, we design a schedule-based package market in which traders submit quantity-contingent bids. Second, we provide a conceptual replication of previous studies of the CDA and avoidable costs. Previous studies had used stationary repetition but we use new cost and value draws each period. Third, we provide experimental evidence that mechanisms that allow traders to express non-convex supply substantially restore efficiency in environments where standard bilateral institutions perform poorly.

In the laboratory, the schedule mechanisms substantially outperform the two CDA variants in both efficiency and in preventing seller losses. In our data, average efficiency rises from 59 percent in CDA Hidden to 61 percent under CDA Visible to 78 percent under One-shot Schedules and to 87 percent under Revisable Schedules, while the incidence of periods with seller losses falls sharply under the schedule mechanisms. These performance gains are robust across environments in which competitive prices do not exist, the core may be empty, or implementing the efficient allocation requires price discrimination.

Our findings suggest a practical way to mitigate market failures that can arise in the presence of avoidable costs. In electricity markets, for example, there are significant "ramping" costs. Once a generator operates at a certain level it is costly to change, which may cause overproduction and negative prices when other sources (e.g. solar panels) produce more electricity than anticipated. To deal with these non-convexities, producers should be allowed to submit "all or nothing" bids in the form of a schedule, i.e. a list of quantity-price pairs with the understanding that if a price from this list is selected then the market operator buys the full associated quantity. Schedules are natural since electricity is a homogeneous commodity and they could be implemented via a "day ahead" market.

1.1. Related Literature

Our work draws on previous research on the CDA, avoidable cost structures, and market design.

Vernon Smith's (1962) seminal study established that the CDA can produce efficient outcomes even with few traders who do not have information about others' costs or values. Smith's experiments used stationary repetition, that is, participants traded in repeated markets with the same costs and values. Smith studied an environment in which costs were convex and values were concave. This meant that a seller who sold at prices above their cost or a buyer who bought at prices below their value would never make a loss from trading.

Some later studies considered environments without stationary repetition. For example, Cason and Friedman (1996) ran CDAs with costs and values randomly redrawn each round. They observed lower efficiency than under stationary repetition. Davis and Williams (1997) studied call markets and CDAs with periodic demand shocks. They observed that prices in call markets adapt to demand shocks more quickly than in comparable CDAs.

Mestelman and Welland (1988) examined experimental markets with advanced production. They found that introducing sunk costs can lead to suboptimal production decisions, reducing market efficiency relative to the case in which production occurs only when an item is sold. Avoidable costs differ from sunk costs in that they are only incurred if a quantity greater than zero is produced.

Van Boening and Wilcox (1996) conducted seminal experiments on avoidable costs, showing that the CDA, normally the most reliably efficient institution, produces erratic "roller coaster" efficiency dynamics, driven by sellers whose participation is necessary for efficiency but who repeatedly incur losses. Subsequent studies investigated whether redesigned institutions could do better. Durham et al. (1996) found that a centralized "smart market" with two-part pricing outperformed the CDA on average but could not sustain full efficiency. Van Boening and Wilcox (2005) showed that the CDA's failure extends to other bilateral contracting institutions. Archibald et al. (2002) found that even explicit cooperation among sellers, with communication and side payments, achieved efficiency only under unusually favorable procedural conditions. Durham et al. (2004) showed that in posted-offer markets loss-making sellers sustain prices above competitive levels through tacit price signaling. Elmaghraby and Larson (2012) found that in procurement auctions with avoidable fixed costs, bidding deviated persistently from equilibrium, reflecting the strategic complexity of the setting. Surveying this research program, Van Boening and Wilcox (2008) concluded that no competitive institution had yet been found that consistently delivers high efficiency under avoidable costs.

Our work also builds on the literature on package mechanisms. The main focus of this literature has been auctions where buyers view items as complements. Without package bidding, they face an exposure problem when there is a risk of acquiring some but not all of the items they want. Early experimental work on package auctions was motivated by a range of potential practical applications. Rassenti et al. (1982) designed and tested auctions for buying airport runway time slots, Banks et al. (1989) introduced and tested the AUSM package auction mechanism motivated by the problem of allocating Space Station resources. Some features of the AUSM mechanism were later incorporated into the RAD mechanism of Kwasnica et al. (2005). Brewer and Plott (1996) designed and tested a combinatorial auction for rights to use railway tracks. In the late 1990s and early 2000s, spectrum auctions emerged as a key practical application of package bidding. Milgrom (2007) provides an overview and many of the key papers can be found in the collection edited by Bichler and Goeree, eds (2017). Another strand of the literature has considered reverse auctions, for example electricity markets with a single buyer and multiple sellers with non-convex costs, see Liberopoulos and Andrianesis (2016) for a review.

Package bidding in markets has been studied less than in auctions. An early example of a package market is the RECLAIM emissions permit scheme. Ishikida et al. (2001) tested the market design using experiments and Fine et al. (2017) describe the details of the mechanism. More recently, the FCC used the 'incentive auction' (Milgrom and Segal, 2020) to reallocate radio spectrum in the US. Bichler et al. (2019) designed and built a combinatorial exchange for the reallocation of catch shares in the New South Wales fishing industry. Milgrom and Watt (2026) drew "inspiration from the novel market design for fisheries by Bichler et al. (2019)" to develop an improved markup mechanism for which they bound the welfare loss. Woods et al. (2025) design and test a continuous-time financial exchange in which traders can submit orders that combine several assets.

Finally, our previous work on package markets includes an exploration of general equilibrium stability in markets with income effects (Goeree and Lindsay, 2016). We found that the unstable prices that arise in the CDA can be tamed by using a mechanism in which traders submit demand and supply schedules. We also examined the exposure problem in markets with differentiated goods (Goeree and Lindsay, 2020) and found that the inefficient outcomes produced by the CDA can be eliminated by allowing package bidding.

1.2. Organization

The remainder of the paper is organized as follows. Section 2 describes the environment and illustrates the phenomena induced by avoidable costs. Section 3 presents the mechanisms and our hypotheses. Section 4 describes the experimental design, and Section 5 reports the results. Finally, Section 6 concludes. An Appendix contains additional outcome figures. Screenshots of the software, examples of submitted schedules from the two schedule treatments (including an animated illustration of the Revisable Schedules mechanism), sample experimental instructions, and control questions are provided in the Supplementary Materials.

2. Environment

We consider an environment where agents have quasi-linear utility and there is a single good that comes in discrete units. Sellers can supply any number of units of the good up to their capacity. Sellers have avoidable costs, which are incurred if they provide one or more units. Sellers have no other costs; hence, the marginal cost of the first unit is the avoidable cost and the marginal cost of subsequent units is zero. Avoidable costs are equivalent to non-sunk fixed costs. They differ from sunk fixed costs in that they are not incurred if output is zero. Buyers can buy any number of units up to their capacity.

Solving the following gives an optimal allocation.

$$\begin{aligned}\max_{\boldsymbol{q}} \sum_{b\in B} v_b(q_b) - \sum_{s\in S} c_s(q_s)\\ \text{subject to} \sum_{b\in B} q_b = \sum_{s\in S} q_s\end{aligned}$$

Note that for a solution to the above, it will always be possible to find a set of cash transfers such that implementing the optimal allocation of the good together with the transfers makes no one worse off. For example, paying each seller the cost of what they supply and charging buyers their value for what they buy.

The presence of avoidable costs can lead to a range of interesting economic phenomena. Table 1 shows some examples.

Existence of competitive prices Let $\mathcal{D}_b(p)$ represent buyer b's demand correspondence and $\mathcal{S}_s(p)$ represent seller s's supply correspondence. A feasible allocation $\boldsymbol{q}$ and price p constitute a competitive equilibrium (CE) if and only if $q_b \in \mathcal{D}_b(p)$ for all $b \in B$ and $q_s \in \mathcal{S}_s(p)$ for all $s \in S$. An example of a market where a CE exists is shown in column 'CE' of Table 1. Seller S1 has an average cost of 3.5 and seller S2 has an average cost of 4. Hence, at prices between 3.5 and 4, market supply is 2. At prices between 3.5 and 4, buyer 1 demands zero units, while buyers 2 and 3 each demand one unit, making market demand equal to 2. Hence, a competitive equilibrium exists.

Uniform price and rationing The optimal allocation can be supported by a single price without making any trader worse off. However, at this price supply and demand are not equal. Column 'No CE' in Table 1 illustrates such a market. The scenario was created by taking 'CE' and increasing S1's cost to 10, making the seller's average cost 5. The efficient outcome is for S1 to sell one unit to B2 and one unit to B3, giving a surplus of 2. However, S2 has a lower average cost than S1, so this outcome cannot be supported as a competitive equilibrium.

Table 1: Scenarios with avoidable costs

	CE	No CE	PD	Empty core
S1, supplies 2	7	<u>10</u>	7	7
S2, supplies 3	12	12	12	12
B1, demands 1	1	1	1	<u>6</u>
B2, demands 1	6	6	<u>2</u>	6
B3, demands 1	6	6	6	6
CE prices	[3.5,4]	∅	∅	∅
linear prices	Yes	Yes	No	Yes
non-empty core	Yes	Yes	Yes	No

Notes: The table shows four scenarios, headed 'CE' (competitive prices exist), 'No CE' (competitive prices do not exist), 'PD' (price discrimination is needed), and 'Empty core', that illustrate markets with different features. The columns show sellers' costs and buyers' values. For example, in CE, seller S1 can supply 2 units for a total cost of 7; buyer B2 values 1 unit at 6. Notice that the 'No CE', 'PD', and 'Empty core' columns each differ from the 'CE' column by one entry, which is underlined.

Price discrimination To achieve the maximum surplus without making any trader worse off than if they had not traded, some traders must trade at different prices. Column 'PD' in Table 1 illustrates such a market. The scenario was created by taking 'CE' and decreasing B2's value to 2. The efficient outcome is still for seller S1 to sell to B2 and B3. However, B2's value is below S1's average cost, hence to cover S1's cost without charging buyers more than their value, B2 and B3 must be charged different prices.

Empty/non-empty core Column 'Empty core' in Table 1 illustrates a market with an empty core. The scenario was created by taking 'CE' and increasing B1's value to 6. Seller S2 can sell to all three buyers and generate a surplus of 6. Seller S1 can sell to two of the buyers and generate a lower surplus of 5. However, there is no way to split the surplus of 6 between S2 and the three buyers without there being a blocking coalition of S1 and two buyers that can do better. Since S1 and any two buyers can achieve a surplus of 5 on their own, avoiding a blocking coalition requires $u_{B_i}+u_{B_j} \geq 5$

for every pair of buyers $i \neq j$, where u_{B_i} denotes buyer B_i's surplus. Summing these three inequalities gives $2(u_{B_1} + u_{B_2} + u_{B_3}) \geq 15$, so the buyers collectively require at least 7.5. This exceeds the total surplus of 6, hence the core is empty.

3. Market Mechanisms

This section describes the four trading mechanisms we evaluate: two variants of the CDA and two variants of the schedule-package market.

3.1. Continuous Double Auction

In both variants of the CDA, traders submit limit orders in continuous time and trade occurs instantly when a pair of compatible orders has accumulated. Orders are price-quantity pairs. Partial filling of orders is allowed. The two variants, CDA Hidden and CDA Visible, differ in whether traders can observe the contents of the order book. Varying the order book visibility lets us test whether sellers can use the information about demand revealed by the order book to avoid falling prey to exposure.[1]

3.2. Schedule-Package Market

In the schedule-package market, traders submit quantity-contingent cash bids. For each quantity, a trader specifies the total cash transfer at which they are willing to buy or sell that quantity. One can think of this as sellers submitting a cost function and buyers submitting a utility or value function. The first variant of the mechanism,

[1] In other contexts, several experimental studies have varied either the information available from the order book or how it is presented. Bloomfield et al. (2015) allow traders to hide some or all of an individual order. Hendershott et al. (2022) allow traders to route orders to either a 'lit' or a 'dark' venue. Ikica et al. (2023) suppress the entire order book and information about others' transactions in their Black Box treatment. Williams (2022) employs a 2×2 design in which only the best bid and ask or all bids and asks, and only own transactions or all transactions, are shown. Finally, Friedman et al. (2026) hold the information set fixed while presenting the order book textually or graphically.

One-shot Schedules, uses a sealed bid format. Once all traders have submitted their schedules, winning bids and payments are determined. The second variant, Revisable Schedules, uses a progressive approach. As schedules are submitted, traders observe the provisional transactions and can update their schedules if desired. The revisions can be submitted in continuous time (rather than discrete rounds). A revision must be an improvement on the previously submitted schedule. The revised schedule must be weakly better at every quantity and strictly better at some quantity. Once there have been no new submissions that have changed the allocation of goods for a period of time, the standing provisional transactions are implemented. The motivation for the 'soft close' was to reduce the incentive for late bidding (see Roth and Ockenfels (2002) for a study of the effects of auction ending rules). In both variants, winning bids and payments are determined using the same mechanism, which is described next.

Winner determination Winning bids are determined as follows. Let B and S denote the sets of buyers and sellers who have submitted a schedule, and let $I = B \cup S$ denote the set of all such traders. For each trader $i \in I$, let $Q_i = \{0, 1, \ldots, q_i^{\max}\}$ denote the quantities they are bidding on, where $q_i^{\max} \in \mathbb{Z}_{>0}$. Each buyer $b \in B$ reports a value schedule $\hat{v}_b : Q_b \to \mathbb{R}$ with $\hat{v}_b(0) = 0$, where $\hat{v}_b(q)$ is the most buyer b offers to pay for q units. Each seller $s \in S$ reports a cost schedule $\hat{c}_s : Q_s \to \mathbb{R}$ with $\hat{c}_s(0) = 0$, where $\hat{c}_s(q)$ is the least seller s asks to receive for supplying q units. For $q \geq 1$, there are no restrictions on the reported values $\hat{v}_b(q)$ and costs $\hat{c}_s(q)$. The mechanism selects a quantity q_i for each trader $i \in I$, with $q_i = 0$ representing not trading. The traded quantities $\boldsymbol{q} = (q_i)_{i \in I}$ are a solution to the following integer optimization problem.[2,3]

[2] The winner determination problem can be solved as an integer linear program by introducing indicator variables for each level of each q_i.

[3] The winner determination problem may have multiple solutions. The mechanism does not require a particular tie breaking rule to be used. In the experimental implementation, ties were broken by selecting the allocations with the greatest sum of winning buyer bids, that is $\sum_{b \in B} \hat{v}_b(q_b)$.

$$
\begin{aligned}
\max_{\boldsymbol{q}} \quad & \sum_{b \in B} \hat{v}_b(q_b) - \sum_{s \in S} \hat{c}_s(q_s) && \text{(maximize reported surplus)} \\
\text{subject to} \quad & \sum_{b \in B} q_b = \sum_{s \in S} q_s && \text{(feasible without disposal)} \\
& q_i \in Q_i \quad \text{for all } i \in I.
\end{aligned}
$$

This is the planner's problem of Section 2 with the reported schedules in place of the traders' true values and costs. The equality constraint means that the total quantity supplied by sellers must equal the total quantity purchased by buyers. In principle, we could allow free disposal, but whether this is appropriate would depend on the application (an airline can fly with empty seats but excess electricity generation can damage grid infrastructure). Let $\boldsymbol{q}^*$ denote a solution. The solution determines the winning bids: trader i's bid at quantity q_i^* is winning whenever $q_i^* > 0$; if $q_i^* = 0$, trader i does not trade and none of their bids is winning.

Pricing rule The mechanism is strictly budget balanced. If traders with winning bids simply paid the amount they bid, the mechanism would generally create a cash surplus. When trade occurs, to distribute this surplus among traders, the mechanism computes prices that are as close as possible to market-clearing prices. It uses an approach similar to Kwasnica et al.'s (2005) RAD mechanism and Fine et al.'s (2017) ACE mechanism. The winning quantities $\boldsymbol{q}^*$ are fixed. Let $W = \{\, i \in I : q_i^* > 0 \,\}$ denote the set of winning traders, with winning buyers $W_B = W \cap B$ and winning sellers $W_S = W \cap S$. A per unit price p and 'distortion' ϵ are found that solve the following.

This meant zero-surplus trade was favored over no trade. Any remaining ties were resolved by the solver's default selection.

$$
\begin{aligned}
\min_{\epsilon,\, p} \quad & \epsilon && && \text{(minimize max distortion)} \\
\text{subject to} \quad & \epsilon \geq 0 \\
& p \geq 0 \\
& \hat{v}_b(q_b^*) - pq_b^* + \epsilon \geq 0 && \text{for all } b \in W_B && \text{(winning buyers)} \\
& pq_s^* - \hat{c}_s(q_s^*) + \epsilon \geq 0 && \text{for all } s \in W_S && \text{(winning sellers)} \\
& \hat{v}_b(q) - pq - \epsilon \leq 0 && \text{for all } b \in B \setminus W,\ q \in Q_b && \text{(losing buyers)} \\
& pq - \hat{c}_s(q) - \epsilon \leq 0 && \text{for all } s \in S \setminus W,\ q \in Q_s. && \text{(losing sellers)}
\end{aligned}
$$

This step approximates a market clearing price. At unit price p, buyer b's reported surplus from buying q units is $\hat{v}_b(q) - pq$, and seller s's reported surplus from selling q units is $pq - \hat{c}_s(q)$. In a competitive equilibrium, the reported surplus would be non-negative for all winning bids, and non-positive for all bids submitted by losing bidders. The constraints 'winning buyers' and 'winning sellers' apply to the winning bids. The constraints 'losing buyers' and 'losing sellers' apply to the bids submitted by traders with no winning bids. (The $q = 0$ constraints hold trivially, since the reported surplus at $q = 0$ is zero.) The losing bids of winning traders are not included. This prevents a trader's losing bid from influencing how much they pay/receive for a winning bid.[4]

When a solution to the problem above yields $\epsilon^* > 0$, then p^* is uniquely determined. Furthermore, a competitive price does not exist. If all winning bidders traded at the chosen price p^*, some would trade on terms worse than they bid. To avoid this while maintaining strict budget balance, a system of 'taxes' and 'subsidies' is

[4]For example, suppose a buyer is willing to pay 10 for one unit, while a seller asks 6 for one unit and 6 in total for two units. The seller's losing two-unit ask should not reduce the payment they receive when their one-unit ask wins. If it were desired to have winners' losing bids influence the price, the following constraints could be added to the linear program:

$$
\begin{aligned}
& \hat{v}_b(q_b^*) - pq_b^* + \epsilon \geq \hat{v}_b(q) - pq && \text{for all } b \in W_B,\ q \in Q_b \setminus \{q_b^*\}, \\
& pq_s^* - \hat{c}_s(q_s^*) + \epsilon \geq pq - \hat{c}_s(q) && \text{for all } s \in W_S,\ q \in Q_s \setminus \{q_s^*\}.
\end{aligned}
$$

Note that including $q = 0$ in these constraints reproduces the 'winning buyers' and 'winning sellers' constraints, since the reported surplus at $q = 0$ is zero.

implemented as follows. Let

$$
\begin{aligned}
g_b &= \hat{v}_b(q_b^*) - p^* q_b^* && \text{for each } b \in W_B, \\
g_s &= p^* q_s^* - \hat{c}_s(q_s^*) && \text{for each } s \in W_S,
\end{aligned}
$$

so that g_i denotes winning trader i's gross reported surplus at price p^*. Traders with $g_i < 0$ receive subsidies $\sigma_i = \max\{0, -g_i\}$, so that traders with $g_i < 0$ have a surplus of zero after the subsidy. These subsidies are financed by taxes on winning traders with $g_i > 0$. Specifically, choose a threshold $g^\dagger \geq 0$ such that

$$
\sum_{\substack{i \in W \\ g_i > 0}} \max\{0,\, g_i - g^\dagger\} = \sum_{\substack{i \in W \\ g_i < 0}} (-g_i).
$$

Then set $\tau_i = \max\{0,\, g_i - g^\dagger\}$. Final surplus is $g_i + \sigma_i - \tau_i$, and the mechanism is strictly budget balanced.

When a solution yields $\epsilon^* = 0$, then p^* may not be uniquely determined.[5] To resolve this indeterminacy, we select the price that maximizes the minimum surplus of the winners, by solving:

$$
\begin{aligned}
\max_{\pi,\, p} \quad & \pi \\
\text{subject to} \quad & \pi \geq 0 \\
& p \geq 0 \\
& \hat{v}_b(q_b^*) - p q_b^* \geq \pi \quad \text{for all } b \in W_B && \text{(winning buyers)} \\
& p q_s^* - \hat{c}_s(q_s^*) \geq \pi \quad \text{for all } s \in W_S && \text{(winning sellers)} \\
& \hat{v}_b(q) - pq \leq 0 \quad \text{for all } b \in B \setminus W,\ q \in Q_b && \text{(losing buyers)} \\
& pq - \hat{c}_s(q) \leq 0 \quad \text{for all } s \in S \setminus W,\ q \in Q_s. && \text{(losing sellers)}
\end{aligned}
$$

Solving this problem will yield a unique price p^*. All traders with winning bids trade at p^* and the mechanism is strictly budget balanced.

[5] When a competitive price exists, $\epsilon = 0$. The converse need not hold. When $\epsilon = 0$ the outcome may not be a CE since CE also requires that the selected quantity for each of the winners maximize their reported surplus given the price.

Table 2: Examples of orders and outcomes

	Orders	Tax/subsidy	payment	surplus
Example 1: CE exists, $p^* = 18$, $\epsilon = 0$, $\pi = 2$				
Buyer 1	$\underline{\mathbf{[-20, 1]}}$	0	-18	2
Buyer 2	[-18, 1]			
Seller 1	$\underline{\mathbf{[10, -1]}}$	0	18	8
Example 2: CE does not exist, $p^* = 5$, $\epsilon = 1$				
Buyer 1	$\underline{\mathbf{[-4, 1]}}$	1	-4	0
Buyer 2	$\underline{\mathbf{[-9, 1]}}$	-2	-7	2
Seller 1	$[11, -1]$, $\underline{\mathbf{[11, -2]}}$	1	11	0

Notes: Orders are pairs specifying amounts of cash and quantity of the good. Negative values indicate giving and positive ones taking. Winning orders are underlined. For each example, the price produced by the mechanism p^* and the value of ϵ are shown. In Example 1, where $\epsilon = 0$, the minimum surplus among the winning bids π is also shown. In the tax/subsidy column, positive entries denote subsidies and negative entries denote taxes.

Table 2 shows two examples of how the prices and payments are selected. In the first example, any price $p \in [18, 20]$ supports the allocation that maximizes reported surplus. Among the candidate prices, $p^* = 18$ maximizes the minimum surplus, which in this case is Buyer 1's. In the second example, there is no uniform price that supports the optimal allocation. The seller supplies zero for $p < 5.5$, supplies two units for $p > 5.5$, and is indifferent between zero and two units at $p = 5.5$. Buyer 1 is only willing to buy at prices of 4 or below and buyer 2 is willing to buy at prices of 9 or below. Hence, there is no price where supply equals demand. It is efficient for the seller to serve both buyers since doing so gives a net surplus of 2. The mechanism produces a price of 5. At this price, buyer 1 and the seller require a subsidy of 1 to trade. The cost of the subsidies is borne by buyer 2 who has a gross surplus of 4 at price 5. Subtracting a tax of 2 from buyer 2 means that buyer 1 and the seller pay/receive what they bid and buyer 2 retains a surplus of 2. Effectively, the mechanism uses price discrimination to cover the seller's costs.

3.3. Hypotheses

We have three hypotheses about the relative performance of the mechanisms we have presented. Together they imply that order-book visibility, package bidding, and allowing revisions each raise efficiency, so the four mechanisms are ranked from CDA Hidden at the bottom to Revisable Schedules at the top.

Hypothesis H1 (Package Mechanism) *The schedule-package mechanisms achieve higher efficiency and fewer seller losses than the CDA in markets where sellers have avoidable costs.*

Rationale. The CDA performs poorly in markets with avoidable costs because trade is bilateral and sequential: a seller commits the entire avoidable cost with the first unit, before knowing whether sufficient demand will materialize to cover it (Van Boening and Wilcox, 1996). The schedule-package market removes this exposure problem. Because bids are quantity-contingent and cleared centrally, a seller whose asks cover their costs incurs the avoidable cost only as part of an allocation whose payments cover the ask, and so can never trade at a loss. Allowing traders to express their non-convex costs directly should therefore raise efficiency and largely eliminate seller losses.

Hypothesis H2 (Revealing Demand) *The CDA Visible mechanism achieves higher efficiency than the CDA Hidden mechanism.*

Rationale. Part of the CDA's poor performance in markets with avoidable costs may be informational. A seller deciding whether to sell their first unit must forecast the revenue from selling subsequent units; with a hidden order book, this forecast rests on little information, so sellers may sell their first unit when there is not sufficient demand or fail to sell it when there is. Making the order book visible lets sellers condition production on observed aggregate demand, reducing both types of error.

Visibility does not, however, remove the exposure problem itself since trades remain bilateral.

Hypothesis H3 (Revisable Schedules) *The Revisable Schedules mechanism achieves higher efficiency than the One-shot Schedules mechanism.*

Rationale. The package market is not strategy-proof; indeed, in environments like ours with two-sided private information, no individually rational, budget-balanced mechanism can guarantee ex-post efficiency (Myerson and Satterthwaite, 1983), so full efficiency is not the relevant benchmark for any of the mechanisms we study. Even in settings where there is a single efficient way to allocate goods, there will typically be many ways to divide the gains from trade among the traders, the optimal bid for a trader with a given value will usually depend on the strategies of other traders, and the mechanism admits multiple equilibria. There is therefore scope for miscoordination; for example, buyers shading their bids or sellers marking up their costs can prevent an efficient trade. Giving traders the opportunity to revise bids after observing provisional outcomes lets them discover the residual supply and demand they face, so we expect iteration to reduce efficiency losses from miscoordination—through adjustment toward best responses rather than through truthful revelation.[6]

4. Experimental Design

We designed an experiment to test our hypotheses and explore how market design affects outcomes in settings with avoidable costs. Groups consisted of 7 traders with

[6]By the revelation principle, any equilibrium outcome of the revisable mechanism could in principle be replicated by a direct mechanism. However, such a direct mechanism would depend on fine details of the environment, such as the distributions of costs and values and the equilibrium being played (Wilson, 1987), whereas the mechanism with revisable schedules is detail-free. Moreover, with multiple equilibria, the revelation principle is silent on which equilibrium traders coordinate on; our hypothesis is that iteration serves as a coordination device, not that it expands the set of attainable equilibrium outcomes. Consistent with this, static and dynamic formats that are theoretically equivalent are known to perform differently in the laboratory (Kagel et al., 1987).

Table 3: Distribution of costs and values

	Capacity	Cost or value		
		First unit	Remaining units	Average at capacity
Seller 1	8	[160,320]	0	[20,40]
Seller 2	5	[125,225]	0	[25,45]
Seller 3	2	[60,100]	0	[30,50]
Buyers 1-4	3	[1,75]		

Notes: Buyers' per-unit values were integers drawn uniformly from the interval $[1, 75]$. Sellers' avoidable costs were restricted to multiples of the seller's capacity, with all multiples within the interval equally likely (e.g. for Seller 1, possible costs were 160, 168, . . . , 320). This restriction ensured that average per-unit costs were always integers.

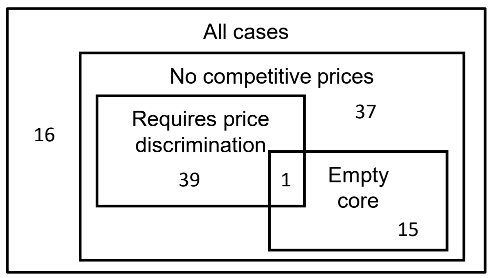

Figure 1: The avoidable cost structures used in the experiment

3 sellers and 4 buyers. Traders had randomly drawn costs and values. Table 3 describes how their costs and values were distributed. A total of 108 cases drawn from the distributions described in the table were used in the experiment. They had a range of interesting features that are summarized in Figure 1.

Our experiment had four treatments, one for each mechanism (screenshots are shown in the Supplementary Materials). In the CDA treatments, periods lasted 4 minutes. Traders submitted limit orders. A transaction occurred as soon as a pair of compatible orders arrived. Submitted orders that had not transacted could be edited or withdrawn. There were two variants: CDA Hidden, in which traders could not see other traders' orders, and CDA Visible, in which all orders in the order book were visible. In both variants, completed transactions were visible. In the schedule treatments, buyers' schedules reported values for a range of quantities, and sellers' schedules reported costs for providing a range of quantities. In the One-shot Schedules treatment, once everyone had submitted a schedule or 4 minutes had passed, the optimal allocation based on the submitted schedules was found and implemented. In the Revisable Schedules treatment, after each schedule was submitted, the provisional transactions were determined and shown. While the market was open, traders could update their schedules, subject to an improvement rule. The market was open for at least 3 minutes, with a soft close to deter sniping: if a schedule submitted at time θ changed the provisional allocation (i.e. the quantities traded), the closing time was updated to $\max\{T, \theta + 30\}$, where T is the planned closing time prior to the submission and time is measured in seconds. Hence, every change to the provisional allocation left all traders at least 30 seconds to react, while submissions that did not alter the provisional allocation never extended the trading time.

There were nine groups of 7 traders per treatment, giving a total of 252 subjects. Instructions were provided in two parts: a slide presentation, which included screenshots of the software and explained how the user interface worked (sample slides are reproduced in the Supplementary Materials), and a paper handout focusing on how payoffs were determined (also in the Supplementary Materials). Participants then completed a set of control questions testing their comprehension of the instructions (Supplementary Materials) and three practice periods before proceeding to the paid part of the experiment. There were 12 paid periods. Groups were fixed, as were buyer and seller roles. Traders received new value and cost draws each period, rather than stationary repetition. Draws were matched across treatments, so, for instance, seller 1 in group 1 and period 3 had the same avoidable cost in each treatment. Seller capacities were rotated each period.

The experiment was conducted at the University of Zurich. Given the technical complexity of the environment and procedures, we restricted recruitment to students studying at ETH Zurich (where students are typically trained in STEM disciplines). The experiment took 2 to 2.5 hours, and the average payment was 47.90 CHF. Payments were based solely on experimental earnings (there was no separate show-up fee), with total earnings calculated as the sum of the values of participants' end-of-period holdings across the 12 paid periods and converted at a rate of CHF 1 per 60 points.

5. Results

Our first result concerns the effect of the market mechanism on outcomes.

Finding 1 (Market design) *Both schedule mechanisms achieve higher efficiency and fewer seller losses than both CDA mechanisms. Efficiency is highest under REVISABLE SCHEDULES, followed by ONE-SHOT SCHEDULES, while the two CDA mechanisms perform similarly.*

Table 4 reports the efficiency achieved in each treatment, measured as the percentage of the potential gains from trade that are realized. The ordering of the treatment means follows the directional predictions of H1, H2, and H3: efficiency is highest under REVISABLE SCHEDULES, followed by ONE-SHOT SCHEDULES, CDA VISIBLE, and CDA HIDDEN.

For all treatment comparisons, the independent unit of observation is the group. We aggregate each outcome over the 12 paid periods within each group and treatment. Groups are then paired across treatments according to their common sequence of cost and value draws, giving $n = 9$ matched pairs. We use exact one-sided Wilcoxon signed-rank tests.

H1 predicts that the schedule mechanisms achieve both higher efficiency and fewer seller losses than the CDA mechanisms. This gives four directional comparisons for each outcome. Mean efficiency is between 16.9 and 27.9 percentage points higher under the schedule mechanisms, and we reject all four component null hypotheses for efficiency ($p < 0.05$ in every comparison). Seller losses are also substantially less frequent: at least one seller makes a loss in 54.6% of periods under CDA Hidden and 41.7% under CDA Visible, compared with 2.8% under One-shot Schedules and 3.7% under Revisable Schedules. We reject all four component null hypotheses for seller losses ($p = 0.002$ in every comparison). Taken together, these results support H1.[7]

H2 predicts that displaying the order book increases efficiency in the CDA. Mean efficiency is 2.2 percentage points higher under CDA Visible than under CDA Hidden, but the one-sided test does not reject the null hypothesis (H_0: CDA Visible $\leq$ CDA Hidden; $p = 0.285$). We therefore find no statistically significant evidence in support of H2.[8]

H3 predicts that allowing schedules to be revised increases efficiency. Mean efficiency is 8.8 percentage points higher under Revisable Schedules than under One-shot Schedules. The one-sided test rejects (H_0: Revisable Schedules $\leq$ One-shot Schedules; $p = 0.049$), providing evidence in support of H3.

The efficiency gains under the schedule mechanisms are reflected mainly in higher seller payoffs, while buyer payoffs are relatively similar across treatments. This pattern is consistent with sellers, rather than buyers, facing the exposure problem. Buyer losses are rare under all four mechanisms. Plots showing the distributions of group-level outcomes (one observation per group) are in Appendix A, while distributions of

[7] The efficiency tests gave $p = 0.014$ for (H_0: One-shot Schedules $\leq$ CDA Hidden), $p = 0.020$ for (H_0: One-shot Schedules $\leq$ CDA Visible), $p = 0.002$ for (H_0: Revisable Schedules $\leq$ CDA Hidden), and $p = 0.002$ for (H_0: Revisable Schedules $\leq$ CDA Visible). The seller-loss tests gave $p = 0.002$ for each of (H_0: One-shot Schedules $\geq$ CDA Hidden), (H_0: One-shot Schedules $\geq$ CDA Visible), (H_0: Revisable Schedules $\geq$ CDA Hidden), and (H_0: Revisable Schedules $\geq$ CDA Visible). In the first set of inequalities, the treatment labels refer to efficiency; in the second, they refer to seller-loss frequency.

[8] Williams (2022) also finds little effect of "opening the order book," i.e. showing all bids and asks versus showing only the best bid and ask.

Table 4: Efficiency, payoffs, and losses

	CDA Hidden	CDA Visible	One-shot Schedules	Revisable Schedules
Efficiency				
Group	59.1 (4.8)	61.3 (5.9)	78.2 (3.4)	87.0 (2.3)
Payoffs				
Seller 1	-16.7 (9.9)	-4.6 (5.1)	13.5 (1.9)	20.7 (4.3)
Seller 2	-4.2 (2.0)	-4.9 (4.7)	9.4 (1.7)	6.6 (3.0)
Seller 3	0.7 (1.8)	1.9 (1.8)	1.6 (0.4)	1.3 (0.5)
Individual Buyer	28.6 (2.8)	26.6 (2.5)	25.0 (1.9)	27.6 (1.7)
Losses				
Seller 1	36.1 (8.3)	23.1 (3.6)	1.9 (1.2)	0.9 (0.9)
Seller 2	20.4 (3.1)	21.3 (4.0)	0.9 (0.9)	1.9 (1.2)
Seller 3	8.3 (3.4)	10.2 (3.3)	0.0 (0.0)	0.9 (0.9)
Any seller	54.6 (9.8)	41.7 (4.8)	2.8 (2.0)	3.7 (2.0)
Individual Buyer	0.0 (0.0)	0.7 (0.3)	0.7 (0.5)	0.2 (0.2)
Any buyer	0.0 (0.0)	2.8 (1.4)	2.8 (2.0)	0.9 (0.9)

Notes: The efficiency figure reports the mean percentage of the potential gains from trade that are realized. The 'Payoffs' panel shows the average payoff per round for each trader role. The 'Losses' panel reports the frequency of losses as a percentage of rounds. Rows naming a specific seller give the percentage of rounds in which that seller made a loss. 'Individual Buyer' gives the percentage of rounds in which a given buyer made a loss (equivalently, the per-round frequency of losses averaged across the four buyers). In contrast, 'Any seller' and 'Any buyer' are market-level measures: the percentage of rounds in which at least one of the three sellers (respectively, one of the four buyers) made a loss. Standard errors of the means are in parentheses.

individual payoffs (one observation per trader-period) are in Appendix B.

It is natural to wonder whether the superior performance of the schedule mechanisms is general or instead occurs only in environments with specific features. Our second result addresses this question.

Finding 2 (Environment) *The ranking of mechanism performance is constant across environments in which competitive prices do or do not exist, in which the core is empty, and in which obtaining the optimal allocation requires price discrimination.*

Table 5: Environment and efficiency

	CDA Hidden	CDA Visible	One-shot Schedules	Revisable Schedules	% of periods
CE exists	66.9 (13.0)	68.5 (10.4)	75.2 (13.9)	79.9 (14.5)	14.8
no CE	58.1 (4.9)	59.2 (6.9)	79.1 (2.6)	87.4 (2.3)	85.2
PD required	55.6 (8.0)	57.7 (12.0)	75.4 (6.2)	87.8 (2.6)	37.0
empty core	58.0 (7.4)	71.3 (8.0)	81.5 (4.3)	88.6 (7.5)	14.8

Notes: The efficiency figures are the mean percentage of the potential gains from trade that are realized. The standard errors of the means are shown in parentheses. 'CE' stands for competitive equilibrium and 'PD' for price discrimination.

Table 5 shows efficiency by treatment for different environment types. Note that the environment types are not mutually exclusive, as illustrated by Figure 1 in Section 4. The ranking of the mechanisms by efficiency is the same across environment types. This suggests that the superior performance of the schedule mechanisms is not driven by specific features of the environment, such as the non-existence of competitive prices.

5.1. CDA

This section focuses on individual behavior in the two CDA treatments. Table 4 showed that the CDA treatments featured low efficiency and frequent losses. In the avoidable cost environment, sellers face a dilemma when deciding whether to sell their first unit and incur their avoidable cost. The revenue from the first unit will rarely be sufficient to cover the avoidable cost on its own. Hence, the expected payoff from the continuation game matters—that is, the amount of revenue the seller expects to receive from selling subsequent units. Table 6 reports regressions of sellers' revenue among those who entered (traded). The revenue from the first sale and subsequent sales are analyzed separately. Note that entry is endogenous and revenue is an equilibrium outcome that depends on other sellers' entry decisions, so these coefficients describe correlates of revenue within the selected sample and realized market condi-

Table 6: Regression analysis of sellers' revenue

	Seller 1		Seller 2		Seller 3	
	unit 1	units 2-8	unit 1	units 2-5	unit 1	unit 2
Seller 1 cost	0.046***	−0.005	−0.009	−0.023	−0.004	0.015
	(0.015)	(0.100)	(0.011)	(0.105)	(0.016)	(0.034)
Seller 2 cost	0.058**	0.608***	0.104***	−0.032	−0.008	0.083
	(0.027)	(0.164)	(0.025)	(0.167)	(0.038)	(0.081)
Seller 3 cost	0.051	0.150	0.043	0.067	0.346***	0.226
	(0.109)	(0.459)	(0.055)	(0.423)	(0.086)	(0.211)
Sum buyer values	0.051**	0.862***	0.053***	0.376***	0.005	0.071
	(0.021)	(0.129)	(0.017)	(0.128)	(0.034)	(0.052)
CDA Visible	−0.326	11.710	0.837	8.009	2.702	0.320
	(2.593)	(12.834)	(1.726)	(9.863)	(2.348)	(5.324)
Constant	4.538	−94.151	10.836	50.247	17.613**	−12.757
	(11.726)	(64.501)	(8.113)	(41.626)	(6.731)	(18.201)
Observations	105	105	95	95	71	71

Notes: Standard errors clustered at the group level are shown in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

tions, not counterfactual revenues under alternative entry patterns. Consider seller 1. Their revenue from selling the first unit increases with their own cost. This is consistent with higher cost sellers requiring a higher price to trade. In addition, it increases with the costs of the other sellers and the sum of buyer values. This can be interpreted as revenue increasing with residual demand. In contrast, for additional units, there is little or no effect of the seller's own cost; however, the effects of the other sellers' costs and buyer values remain. A natural interpretation is that the marginal cost of additional units is zero but the effect of residual demand is unchanged. Sellers 2 and 3 exhibit a similar pattern, although it is less pronounced.

Our next result concerns the dilemma sellers face.

Finding 3 (First-unit revenue) *In both CDA treatments, high-cost sellers frequently*

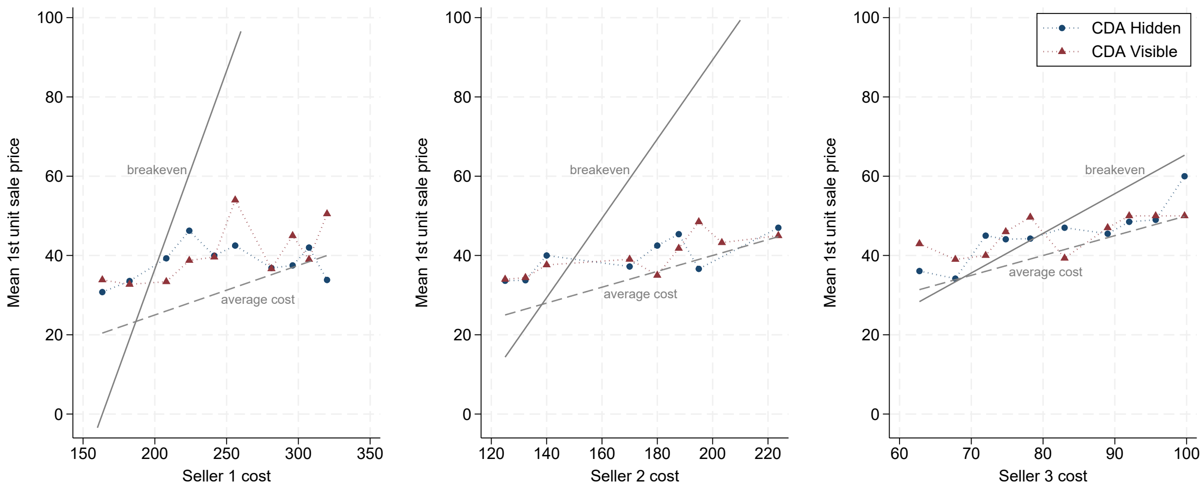

Figure 2: Sellers' revenue from first unit

sell at first-unit prices below the level required to break even given expected continuation-game revenue.

The evidence for this result is drawn from two sources. The regressions reported in Table 6 suggest that the revenue from second and subsequent units is only weakly correlated with sellers' avoidable costs. Figure 2 shows how sellers' first unit revenue varied with cost. For each seller role (seller 1, 2, and 3) and each treatment (CDA Hidden and CDA Visible), a binscatter plot with 10 bins is displayed, showing how revenue varies with seller cost. The linear dashed lines show the seller's average cost. The steeper solid linear lines show the estimated first-unit break-even price, which takes into account the seller's cost and the expected revenue from the continuation game. The observed first-unit revenues are predominantly above the sellers' average cost. However, except for low-cost sellers, first-unit revenues are frequently below the break-even price. This is consistent with high-cost sellers overestimating the continuation game revenue, and as a consequence, mistakenly selling their first unit too cheaply, leading to losses.

The motivation for varying visibility of the order book was that it would affect how much information sellers had about the residual demand, and that this might affect behavior and efficiency. The efficiency results presented in Table 4, the regression analysis in Table 6, and the plots shown in Figure 2 suggest that any effect of order book visibility was relatively small. The next result focuses on the effect of order book visibility.

Finding 4 (Order-book visibility) *While making the order book visible has only a small effect on efficiency, it does change entry behavior by allowing sellers to condition entry on observed aggregate demand.*

We estimate linear probability models of sellers' activity and report the results in Table 7. The coefficients show that the probability of a seller incurring their avoidable cost decreases with the seller's own cost and increases with the costs of other sellers. The effect of buyer values is mixed, with no clear pattern across sellers. This is consistent with negative correlations between different sellers' activity. For seller 1, it appears that with a hidden order book, the maximum value of the buyers matters, whereas with a visible order book, the sum of the buyers' values matters. This can be interpreted as the visible order book allowing Seller 1 to condition activity based on aggregate demand. This could explain why we see slightly higher efficiency and fewer periods with losses (see Table 4) with the visible order book. Sellers knew the distribution from which buyers' values were drawn, but never observed the realized draws, so any correlation between sellers' decisions and buyer values must reflect inference from the orders buyers submitted.

5.2. Schedules

We now turn to the two schedule mechanisms, beginning by analyzing the schedules that traders submitted. In the experiment, subjects were presented with a list of quantities and specified a price per unit for each one. A schedule can be classified as *truthful* if it reports the trader's true cost or value at every quantity. Let t denote

Table 7: Seller Activity

	Seller 1 enters	Seller 2 enters	Seller 3 enters
Seller 1 cost	-0.004^{***}	0.001^{**}	0.001^{**}
	(0.001)	(0.001)	(0.000)
Seller 2 cost	0.001	-0.006^{***}	0.003^{**}
	(0.001)	(0.001)	(0.001)
Seller 3 cost	0.002	0.003	-0.014^{***}
	(0.003)	(0.003)	(0.002)
Max buyer value	0.017^{***}	-0.005	0.004
	(0.004)	(0.004)	(0.004)
CDA Visible	0.109	0.215	-0.133
	(0.196)	(0.268)	(0.223)
CDA Visible × max buyer value	-0.017^{***}	0.004	0.005
	(0.005)	(0.005)	(0.005)
Sum buyer values	-0.002	0.005^{**}	0.002
	(0.001)	(0.002)	(0.001)
CDA Visible × sum buyer values	0.006^{***}	-0.002	-0.001
	(0.002)	(0.002)	(0.001)
Constant	0.336	0.347	0.258
	(0.413)	(0.495)	(0.389)
Observations	216	216	216

Notes: Standard errors clustered at the group level are shown in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

the trader's private type: the avoidable cost for a seller and the per-unit value for a buyer. A truthful seller with avoidable cost t would report the downward-sloping schedule $p = t/q$, and a truthful buyer with per-unit value t would report the flat schedule $p = t$. Schedules can more generally be classified by how the reported price varies with quantity. A *constant-price* schedule specifies the same price per unit at every quantity. A *profit-target* schedule sets prices so that the trader would earn the same profit at every quantity if trades were executed at the reported prices. With target π^{tar}, a seller reports $p = (t + \pi^{\text{tar}})/q$ and a buyer reports $p = t - \pi^{\text{tar}}/q$. Note

Table 8: Analysis of submitted schedules

	One-shot		Revisable, first		Revisable, last	
	buyer	seller	buyer	seller	buyer	seller
Panel A: Shape						
Truthful RMSE	19.48	32.00	54.82	100.25	24.67	39.18
Constant price RMSE	9.23	104.42	7.24	140.47	11.42	119.67
Profit target RMSE	42.76	13.67	20.42	15.59	41.56	21.22
Panel B: Payoffs						
Submitted schedule payoff	25.05	8.15	7.10	1.01	27.59	9.51
Best response payoff	35.44	22.98	33.77	16.34	33.77	16.34
Ex-post regret	10.39	14.83	26.67	15.34	6.19	6.83

that $\pi^{\text{tar}} = 0$ for truthful buyers and sellers.

Finding 5 (Schedule shape) *Submitted schedules are systematically non-truthful: buyers are best described by constant-price schedules, while sellers are best described by profit-target schedules.*

We analyzed how well the submitted schedules fit each of the classifications using root mean squared error (RMSE) as the measure of fit. The results are reported in Panel A of Table 8. Buyers and sellers are considered separately. In the Revisable Schedules treatment, traders could update their submissions. The entries in the 'Revisable, first' column are based on each trader's initial schedule, while those in the 'Revisable, last' column are based on the final schedule. The constant-price and profit-target parameters were estimated separately for each schedule. To guard against overfitting, we used leave-one-out cross-validation: for a schedule with n points, the fit of the j^{th} point was evaluated using the parameter estimated from the other $n-1$ points.

Several patterns emerge. For buyers, the constant-price classification provides the best fit, whereas for sellers the profit-target classification provides the best fit. To assess whether these differences are systematic across independent groups, we first averaged the schedule-level RMSEs within each group and then conducted exact paired Wilcoxon signed-rank tests. In both the One-shot Schedules treatment and the final submissions under Revisable Schedules, constant-price schedules fit buyers better than either truthful or profit-target schedules, while profit-target schedules fit sellers better than either truthful or constant-price schedules. All nine group-level differences have the predicted sign in each comparison. The differences are significant in exact two-sided tests ($p = 0.004$ throughout) and remain significant after applying a Holm correction to the ten comparisons underlying Finding 5 ($p_{\text{Holm}} = 0.039$ throughout). Within the Revisable Schedules treatment, truthful RMSE is lower for final than for initial submissions in all nine groups for both buyers and sellers.

To investigate why the Revisable Schedules mechanism performs well, we compare the payoffs generated by submitted schedules with those available from ex-post best responses.

Finding 6 (Ex-post best responses) *Final submissions in the Revisable Schedules mechanism are closer to ex-post best responses than submissions in the One-shot Schedules mechanism.*

For each trader in each period, we calculated the payoff generated by the submitted schedule and the payoff available from an ex-post best response to the schedules ultimately submitted by the other traders. The best response was found by iterating over the set of candidate schedules and selecting the schedule that generated the highest payoff. As in the experiment, candidate schedules were restricted to integer-valued bids. We define *ex-post regret* as the difference between the best-response payoff and the payoff from the submitted schedule. Panel B of Table 8 reports the mean submitted-schedule payoff, best-response payoff, and ex-post regret.

For the Revisable Schedules treatment, we also evaluate each trader's initial schedule against the other traders' final schedules. Mean regret falls substantially between the initial and final submissions: from 26.67 to 6.19 for buyers and from 15.34 to 6.83 for sellers. Final submissions in Revisable Schedules also have lower mean regret than submissions in One-shot Schedules: 6.19 compared with 10.39 for buyers and 6.83 compared with 14.83 for sellers.

For inference, we averaged regret within each independent group and trader role and paired groups across treatments using their common cost and value draws. Final Revisable Schedules submissions have lower regret than One-shot Schedules submissions in eight of the nine matched groups for both buyers and sellers. Exact two-sided Wilcoxon signed-rank tests reject the null of no difference for each role ($p = 0.020$ for buyers and $p = 0.020$ for sellers). Both comparisons remain significant after applying a Holm correction across the two trader-role comparisons ($p_{\text{Holm}} = 0.039$ in each case). The Supplementary Materials illustrate how a period with the same cost and value draws developed under the two schedule treatments, including an animated version of the Revisable Schedules example.

Our final piece of analysis concerns whether traders best respond ex-ante, that is, against the distribution of strategies played by others. The ex-post benchmark of Finding 6 is a natural measure of how well traders' final schedule submissions are adapted to the final submissions of others. However, it does not account for traders' uncertainty: in One-shot Schedules traders have no information about others' schedules when submitting; in Revisable Schedules traders do not know what the future submissions of other traders will be. The ex-ante analysis can be applied to both treatments and can account for uncertainty. The analysis exploits the random and private assignment of types. Random assignment means that, within each treatment and trader role, observations associated with different type draws are exposed, in expectation, to the same distribution of group- and period-specific strategic environments. Opponents' strategies may differ across groups or evolve over periods; such variation is averaged over rather than assumed absent. Because types are private, opponents can respond to a trader's observed actions but cannot respond directly to whether the trader genuinely has a particular type or is imitating the strategy associated with that type.

We cannot verify that a strategy is an ex-ante best response. In Revisable Schedules, a strategy is a contingent plan, and we observe only the realized path of play: the payoff to an arbitrary alternative strategy cannot be computed because we do not observe how opponents would have behaved at information sets that were not reached. We can, however, identify violations of a necessary condition for best response. A trader of type t could have imitated the strategy associated with another type t' in the same role. Random assignment allows the average allocation and payment observed when type t was drawn—including opponents' responses to that strategy—to be used to estimate the average outcome of this deviation. For example, the strategy associated with sellers of cost t fails this necessary condition if a seller with cost t would have earned a higher average payoff by adopting the strategy associated with cost $t' \neq t$.

For each type draw t of each player role r in each schedule mechanism k, we calculated a pair $(\bar{m}_{trk}, \bar{x}_{trk})$. For a buyer with a private value draw t, $\bar{m}_{trk}$ is the mean expenditure; for a seller with an avoidable cost draw t, $\bar{m}_{trk}$ is the mean revenue. The entry $\bar{x}_{trk}$ is the mean of payoff relevant activity x, which is defined as

$$x = \begin{cases} q, & \text{for a buyer,} \\ \mathbf{1}\{q > 0\}, & \text{for a seller.} \end{cases}$$

Thus, for a buyer, $\bar{x}$ is the mean quantity purchased, whereas for a seller, $\bar{x}$ is the probability of entering and incurring the avoidable cost. The trader's mean payoff can then be written as

$$\bar{\pi} = \begin{cases} t\bar{x} - \bar{m}, & \text{for a buyer,} \\ \bar{m} - t\bar{x}, & \text{for a seller.} \end{cases}$$

For pairs with $\bar{x} > 0$, we can construct an implicit price $\bar{p} = \bar{m}/\bar{x}$. For a buyer, $\bar{p}$ is mean expenditure per unit. For a seller, it is mean revenue conditional on incurring the avoidable cost, rather than a per-unit price.

We now impose additional structure by assuming a linear relationship between $\bar{x}$ and $\bar{p}$. This lets us estimate the implied residual 'supply' and 'demand' curves for each trader role (seller 1, seller 2, seller 3, and buyer) and mechanism using the following model

$$\bar{p}_{trk} = \beta_{0,rk} + \beta_{1,rk}\bar{x}_{trk} + u_{trk}.$$

The linear supply and demand curves plotted in Figure 3 illustrate this for the buyers and Seller 1.[9] From the residual supply curve faced by buyers, the corresponding marginal expenditure curve is $\mathrm{ME}_{rk}(\bar{x}) = \beta_{0,rk} + 2\beta_{1,rk}\bar{x}$. From the residual demand curve faced by sellers, the marginal revenue curve is derived equivalently. The estimated ME and MR curves can be used to predict the outcomes that would occur from playing a payoff-maximizing strategy. Mirroring textbook analysis, a payoff-maximizing buyer should choose a strategy yielding $\bar{x}$ where $\mathrm{ME}(\bar{x}) = t$ and a seller a strategy where $\mathrm{MR}(\bar{x}) = t$. Two caveats distinguish this from the textbook case: $\bar{x}$ and $\bar{m}$ are expected outcomes, with the expectation taken over the distribution of others' strategies, and the fitted curve is an approximation to the menu of $(\bar{m}, \bar{x})$ pairs an individual trader could attain by varying their strategy. The exercise therefore asks whether traders best respond *ex-ante*, against the empirical distribution of others' strategies; this is distinct from the ex-post question, taken up in Finding 6, of how traders' submissions compare with the best response to the schedules actually submitted in a given market. Table 9 reports the difference between actual activity levels and those implied to be optimal by the ME and MR curves. Negative numbers imply not trading enough, positive ones imply trading too readily.

[9] Pairs with $\bar{x} = 0$ are excluded. Note, since non-traders do not pay or receive anything, $\bar{x} = 0$ implies $\bar{m} = 0$.

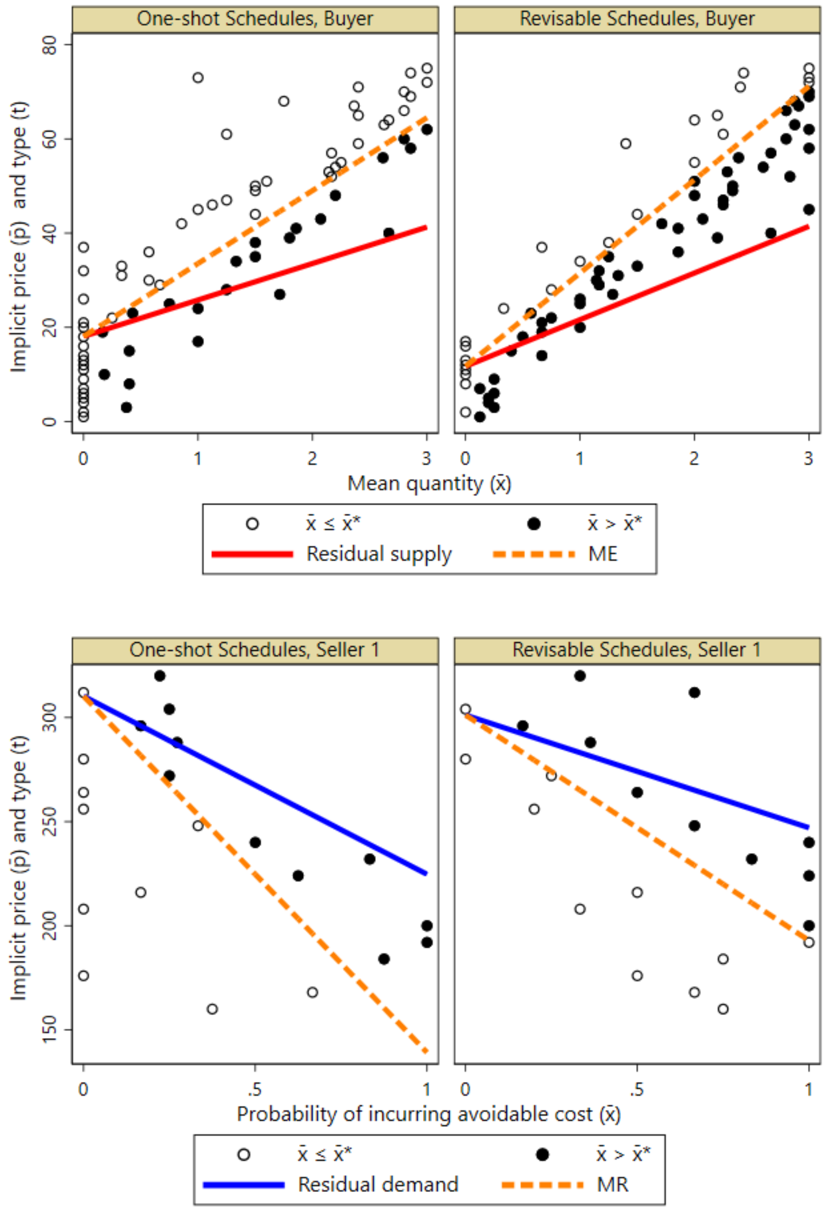

Figure 3: Strategies in One-shot Schedules and Revisable Schedules

Notes: Each point corresponds to a distinct type draw and averages over all markets in which a trader received that draw. For buyers (top), the points plot the buyer's value against the mean quantity traded across markets with that draw. For Seller 1 (bottom), the points plot the seller's avoidable cost against the proportion of markets in which the avoidable cost was incurred. The solid line is the estimated residual supply (demand) curve in terms of the implicit price $\bar{p}$, which for the seller is the average revenue conditional on incurring the avoidable cost rather than a per-unit price; the dashed line is the implied marginal expenditure (marginal revenue) curve. The value $\bar{x}^*$ denotes the $\bar{x}$ at which the marginal curve equals the trader's value (cost).

Table 9: Strategies with schedules

	Buyer	Seller 1	Seller 2	Seller 3
One-shot Schedules	-0.167	-0.055	-0.936	-0.117
	(0.068)	(0.070)	(0.360)	(0.048)
Revisable Schedules	0.154	-0.033	-0.069	-0.124
	(0.045)	(0.078)	(0.057)	(0.076)
Difference (Revisable − One-shot)	0.321	0.022	0.868	-0.006
	(0.072)	(0.060)	(0.335)	(0.062)
p-value (paired t-test)	0.000	0.714	0.020	0.919
p-value (Wilcoxon signed-rank)	0.000	0.683	0.073	0.354

Notes: The statistical tests have one observation per type draw per role.

We observe the following.

Finding 7 (Strategic behavior) *The* Revisable Schedules *mechanism induces buyers to behave more competitively.*

The 'Buyer' column in Table 9 shows a value of -0.167 for the One-shot Schedules treatment and 0.154 for the Revisable Schedules treatment. The null hypothesis of no difference can be rejected (using a paired t-test or a Wilcoxon signed-rank test). The figures suggest that in the One-shot Schedules treatment, buyers could increase profits by bidding slightly higher and trading more frequently, whereas in the Revisable Schedules treatment, the reverse is true. This can be seen in Figure 3, where for the buyer in One-shot Schedules, there are more hollow points (left of the ME curve) whereas in Revisable Schedules, there are more solid points (right of the ME curve). For the sellers, the evidence is mixed and no overall pattern is apparent.

6. Discussion

Markets with avoidable costs pose a problem for standard trading institutions because the efficient allocation is inherently multilateral: a seller can profitably incur its avoidable cost only if enough demand is assembled to cover it, yet bilateral trades commit the seller one unit at a time. The continuous double auction, which performs remarkably well in many environments, inherits this difficulty and produces low efficiency and volatile prices. The central contribution of this paper is to show that this failure is not a feature of the environment but of the trading mechanism. The schedule-package market allows traders to express the non-convex supply that avoidable costs create and restores much of the efficiency that bilateral institutions forgo. In our experiment, average efficiency rises from 59 and 61 percent under the two CDA formats to 78 percent under One-shot Schedules and 87 percent under Revisable Schedules, while the share of periods in which at least one seller makes a loss falls from 54.6 percent and 41.7 percent under the CDAs to 2.8 percent and 3.7 percent under the schedule mechanisms. The mechanisms that allowed traders to express non-convex preferences thus delivered large gains both in efficiency and the avoidance of losses.

The schedule mechanisms eliminate the exposure problem that drives the CDA's poor performance. Under bilateral trading, a seller must sell a first unit before knowing whether subsequent units will follow, and so must gamble on the revenue it expects from the continuation of the market. We saw the consequences of this in the CDA treatments, where high-cost sellers frequently entered at first-unit prices too low to break even given the revenue they could expect to earn from later sales. The schedule-package market eliminates this gamble: a seller that submits asks covering its true avoidable cost cannot incur a loss. Because the pricing rule permits individualized payments—financed by transfers among winning traders while preserving strict budget balance—the mechanism can cover a seller's avoidable cost even when no single uniform price would do so. This is why the advantage of the schedule mechanisms is not confined to any one type of market. Finding 2 showed that the ranking of mechanisms is the same whether or not competitive prices exist, whether the core is empty, and whether achieving the efficient allocation requires price discrimination. The mechanism is not exploiting a favorable special case; it accommodates the full range of phenomena that avoidable costs generate.

The gap between the One-shot Schedules and Revisable Schedules mechanisms points to the role of feedback rather than truthful revelation. Neither mechanism is incentive compatible, and traders did not bid truthfully: buyers' schedules are best described as posting a constant price across quantities, while sellers' schedules are best described by a profit target. Revisions did not undo this strategic behavior, but feedback did let traders learn the residual demand and supply they face and adjust accordingly. Submissions under the Revisable Schedules mechanism move closer to ex-post best responses to the schedules others actually submit. Further, the ex-ante analysis shows that revisions push buyers toward more competitive bidding. Feedback thus helps traders coordinate on an allocation close to the efficient one despite bidding strategically.

When interpreting our results, it is important to recognize that the mechanisms implemented in the laboratory are bundles of design features. The schedule-based mechanisms differ from the CDA in the messages traders submit (quantity-contingent schedules rather than orders for individual units), in when trades occur (only once the submission phase has closed, rather than as soon as a bid and an ask cross), and in how payments are determined (individualized payments financed by transfers among winners, rather than prices set trade by trade). Our goal was not to estimate the separate marginal effect of each feature, but to evaluate whether a coherent market institution could overcome the exposure problem created by avoidable costs.

We therefore compared complete mechanisms that combined features we expected to be complementary: quantity-contingent schedules, centralized package clearing, and, in the Revisable Schedules treatment, opportunities for revision and a soft close. This design does not provide a full factorial decomposition, but it identifies the performance of practically implementable market designs. Future work could isolate individual components; our results show that, as a package, these design features substantially improve efficiency and reduce seller losses.

We close by discussing how the schedule-package market could be applied in practice. We consider three issues: market thickness, the computational cost of winner determination, and the settlement of multilateral trades. First, our lab experiments implemented relatively small markets with seven traders. In thicker markets, the impact of non-convex costs may be less pronounced (see Starr (1969) for an analysis of how the impact of non-convex preferences depends on market size). For example, if airline passengers view alternative flights as perfect substitutes, it is possible to find an efficient allocation of passengers to flights where at most one flight is partially full. However, if passengers have preferences over, say, departure time, an efficient allocation may involve many flights being partially full. This suggests the impact of non-convexity is likely to depend on the specific details of a market, not just its size.

A further consideration is computational complexity. Winner determination in the schedule treatments required solving an integer linear program, and prices were determined by solving linear programs. These optimization problems were tractable in our experiment. In larger applications with multiple goods or other features, however, winner determination may become computationally demanding. Integer linear programming problems are often NP-hard, and the fastest known exact methods for solving them can have worst-case running times that grow rapidly with input size. Some large instances of the schedule-package market may therefore be difficult to solve exactly within the time available for market clearing. This does not mean that the approach is impractical. The computational difficulty of integer programming is a worst-case concern, and many relevant instances can be solved quickly in practice. If computational constraints arise, the mechanism could be adapted in several ways. The winner-determination problem could be solved using heuristics or approximation methods that preserve feasibility but do not guarantee the surplus-maximizing allocation. Alternatively, the market could be decomposed into smaller submarkets, albeit at the cost of potentially losing gains from trade across submarket boundaries. These modifications would trade off some theoretical optimality against computational tractability, but they would preserve the central idea of the design: allowing traders to express quantity-contingent bids so that avoidable costs can be covered through coordinated package trades.

A final consideration concerns settlement. In the package market, transactions often involve more than two traders, which may make defaults more complex to handle than in bilateral markets. If two sellers are matched with three buyers and one seller defaults, the mechanism must specify whether the entire transaction is canceled or whether the remaining supply is allocated among buyers according to some priority rule. Similarly, if a seller is matched with two buyers and one buyer defaults, the payment from the remaining buyer may be insufficient to cover the seller's avoidable cost. Thus, multilateral trades create settlement interdependencies: one participant's default can affect the feasibility or budget balance of trades involving others. Practical implementations would need explicit settlement rules, such as collateral, default penalties, a central counterparty, or re-clearing after a default.

These are practical design choices, not limitations: each of the considerations above – market thickness, the computation of the clearing allocation, and the settlement of multilateral trades – has a workable resolution, in several cases borrowing from the tools already used to clear large non-convex markets such as wholesale electricity. The broader lesson stands independent of these details. Avoidable costs make the efficient allocation multilateral, and bilateral institutions cannot reliably assemble it; a market that lets traders express quantity-contingent bids and clears them as packages can. Avoidable costs are not an insurmountable obstacle to efficient exchange – they are a market-design problem with a market-design solution.

References

Archibald, Glen, Mark V. Van Boening, and Nathaniel T. Wilcox, "Avoidable Cost: Can Collusion succeed where competition fails?," in R. Mark Isaac and Charles A. Holt, eds., *Experiments Investigating Market Power*, Vol. 9 of *Research in Experimental Economics*, JAI Press, Stamford, CT, 2002, pp. 217–242.

Banks, Jeffrey S., John O. Ledyard, and David P. Porter, "Allocating Uncertain and Unresponsive Resources: An Experimental Approach," *The RAND Journal of Economics*, 1989, *20* (1), 1–25.

Bichler, Martin and Jacob K. Goeree, eds, *Handbook of Spectrum Auction Design*, Cambridge University Press, 2017.

_ **, Vladimir Fux, and Jacob K. Goeree**, "Designing Combinatorial Exchanges for the Reallocation of Resource Rights," *Proceedings of the National Academy of Sciences*, 2019, *116* (3), 786–791.

Bloomfield, Robert, Maureen O'Hara, and Gideon Saar, "Hidden Liquidity: Some New Light on Dark Trading," *The Journal of Finance*, 2015, *70* (5), 2227–2274.

Brewer, Paul J. and Charles R. Plott, "A binary conflict ascending price (BICAP) mechanism for the decentralized allocation of the right to use railroad tracks," *International Journal of Industrial Organization*, 1996, *14* (6), 857–886.

Cason, Timothy N. and Daniel Friedman, "Price formation in double auction markets," *Journal of Economic Dynamics and Control*, 1996, *20* (8), 1307–1337.

Davis, Douglas D. and Arlington W. Williams, "The effects of nonstationarities on performance in call markets," *Journal of Economic Behavior & Organization*, 1997, *32* (1), 39–54.

Durham, Yvonne, Kevin McCabe, Mark A Olson, Stephen J Rassenti, and Vernon L Smith, "Oligopoly competition in fixed cost environments," *International Journal of Industrial Organization*, 2004, *22* (2), 147–162.

_ , Stephen J. Rassenti, Vernon L. Smith, Mark V. Van Boening, and Nathaniel T. Wilcox, "Can Core Allocations Be Achieved in Avoidable Fixed Cost Environments Using Two-Part Price Competition?," *Annals of Operations Research*, 1996, *68*, 61–88.

Elmaghraby, Wedad J. and Nathan Larson, "Explaining Deviations from Equilibrium in Auctions with Avoidable Fixed Costs," *Games and Economic Behavior*, September 2012, *76* (1), 131–159.

Fine, Leslie, Jacob K. Goeree, Tak Ishikida, and John O. Ledyard, "ACE: A combinatorial market mechanism," in Martin Bichler and Jacob K. Goeree, eds., *Handbook of Spectrum Auction Design*, Cambridge University Press, 2017, chapter 40, pp. 874–901.

Friedman, Daniel, Brett Williams, and Vivian Juehui Zheng, "Market visualizations," *Experimental Economics*, 2026, *28* (6), 1321–1342.

Goeree, Jacob K. and Luke Lindsay, "Market design and the stability of general equilibrium," *Journal of Economic Theory*, 2016, *165*, 37–68.

_ and _ , "The Exposure Problem and Market Design," *The Review of Economic Studies*, 09 2020, *87* (5), 2230–2255.

Hendershott, Terrence, Marvin Wee, and Yuanji Wen, "Transparency in fragmented markets: Experimental evidence," *Journal of Financial Markets*, 2022, *59*, 100732.

Ikica, Barbara, Simon Jantschgi, Heinrich H. Nax, Diego G. Nuñez Duran, and Bary S. R. Pradelski, "Competitive Market Behavior: Convergence and Asymmetry in the Experimental Double Auction," *International Economic Review*, 2023, *64* (3), 1087–1126.

Ishikida, Takashi, John O. Ledyard, Mark A. Olson, and David P. Porter, "Experimental testbedding of a pollution trading system: Southern California's RECLAIM emissions market," in R. Mark Isaac, ed., *Research in Experimental Economics*, Vol. 8, Emerald Group Publishing Limited, 03 2001, pp. 185–220.

Kagel, John H., Ronald M. Harstad, and Dan Levin, "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study," *Econometrica*, 1987, *55* (6), 1275–1304.

Kwasnica, Anthony M., John O. Ledyard, Dave P. Porter, and Christine DeMartini, "A New and Improved Design for Multiobject Iterative Auctions," *Management Science*, March 2005, *51* (3), 419–434.

Liberopoulos, George and Panagiotis Andrianesis, "Critical Review of Pricing Schemes in Markets with Non-Convex Costs," *Operations Research*, 2016, *64* (1), 17–31.

Mestelman, Stuart and Douglas Welland, "Advance Production in Experimental Markets," *The Review of Economic Studies*, 10 1988, *55* (4), 641–654.

Milgrom, Paul, "Package Auctions and Exchanges," *Econometrica*, 2007, *75* (4), 935–965.

— **and Ilya Segal**, "Clock auctions and radio spectrum reallocation," *Journal of Political Economy*, 2020, *128* (1), 1–31.

— **and Mitchell Watt**, "A Walrasian Mechanism with Markups for Nonconvex Markets," *The Review of Economic Studies*, 05 2026, *93* (3), 1995–2020.

Myerson, Roger B. and Mark A. Satterthwaite, "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 1983, *29* (2), 265–281.

Rassenti, S. J., V. L. Smith, and R. L. Bulfin, "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *The Bell Journal of Economics*, 1982, *13* (2), 402–417.

Roth, Alvin E. and Axel Ockenfels, "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, 2002, *92* (4), 1093–1103.

Scarf, Herbert, "The Allocation of Resources in the Presence of Indivisibilities," *Journal of Economic Perspectives*, December 1994, *8* (4), 111–128.

Smith, Vernon L., "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*, 1962, *70* (2), 111–137.

Starr, Ross M., "Quasi-Equilibria in Markets with Non-Convex Preferences," *Econometrica*, 1969, *37* (1), 25–38.

Telser, Lester G., "The Usefulness of Core Theory in Economics," *Journal of Economic Perspectives*, June 1994, *8* (2), 151–164.

Van Boening, Mark V. and Nathaniel T. Wilcox, "Avoidable Cost: Ride a Double Auction Roller Coaster," *American Economic Review*, 1996, *86* (3), 461–477.

_ **and** _ , "A Limit of Bilateral Contracting Institutions," *Economic Inquiry*, 2005, *43* (4), 840–854.

_ **and** _ , "Avoidable Cost Structures and Competitive Market Institutions," in Charles R. Plott and Vernon L. Smith, eds., *Handbook of Experimental Economics Results*, Vol. 1, Elsevier, 2008, chapter 25, pp. 206–212.

Williams, Brett, "Opening the Book: Price Information's Impact on Market Efficiency in the Lab," 2022. Working Paper.

Wilson, Robert, "Game-theoretic Analyses of Trading Processes," in Truman F. Bewley, ed., *Advances in Economic Theory: Fifth World Congress*, Econometric Society Monographs, Cambridge: Cambridge University Press, 1987, pp. 33–70.

Woods, R., J. O. Ledyard, S. Malick, P. Bossaerts, and N. Yadav, "Continuous Combinatorial Financial Exchange: Design, Simulations, and Live Testing," 2025. Working Paper.

Appendix

A. Distribution of group outcomes

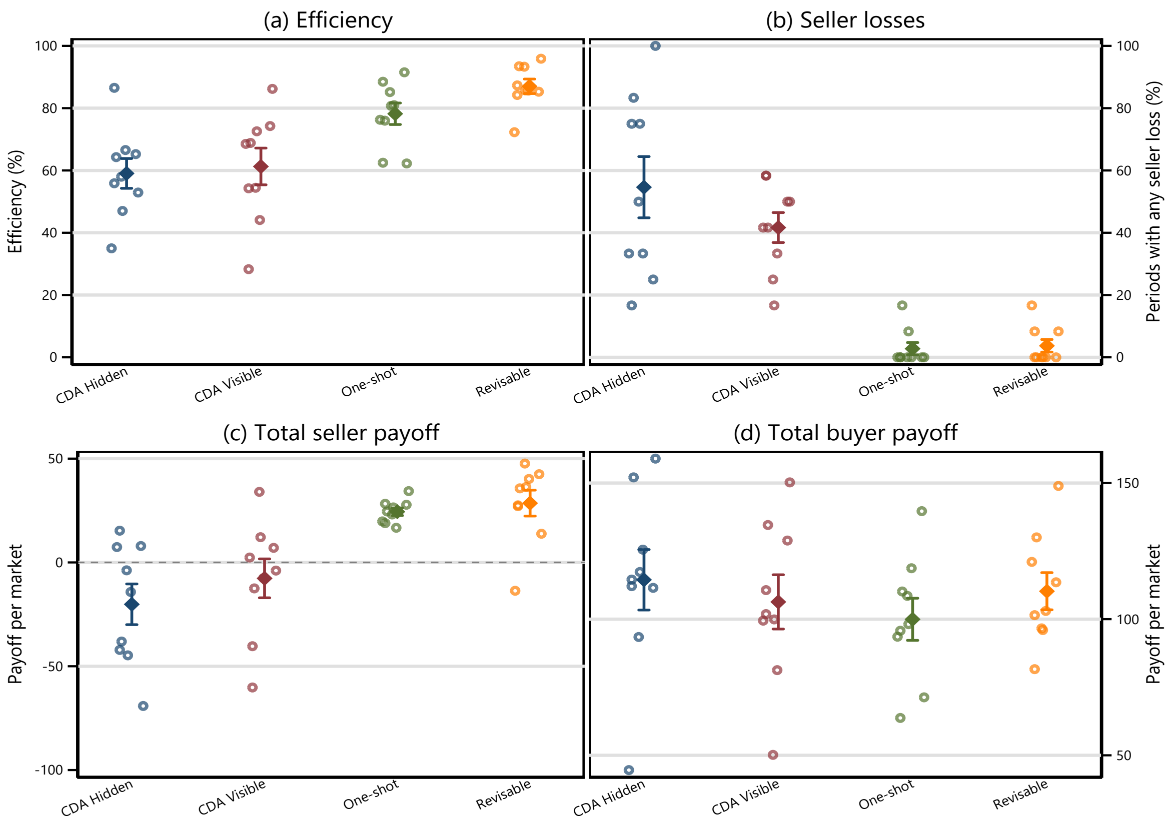

The figure shows group-level outcomes by treatment for (a) efficiency, (b) the percentage of periods with any seller loss, (c) total seller payoff, and (d) total buyer payoff. Open circles show the nine group-level observations in each treatment, diamonds show treatment means, and capped bars indicate plus or minus one standard error of the mean.

B. Distribution of payoffs

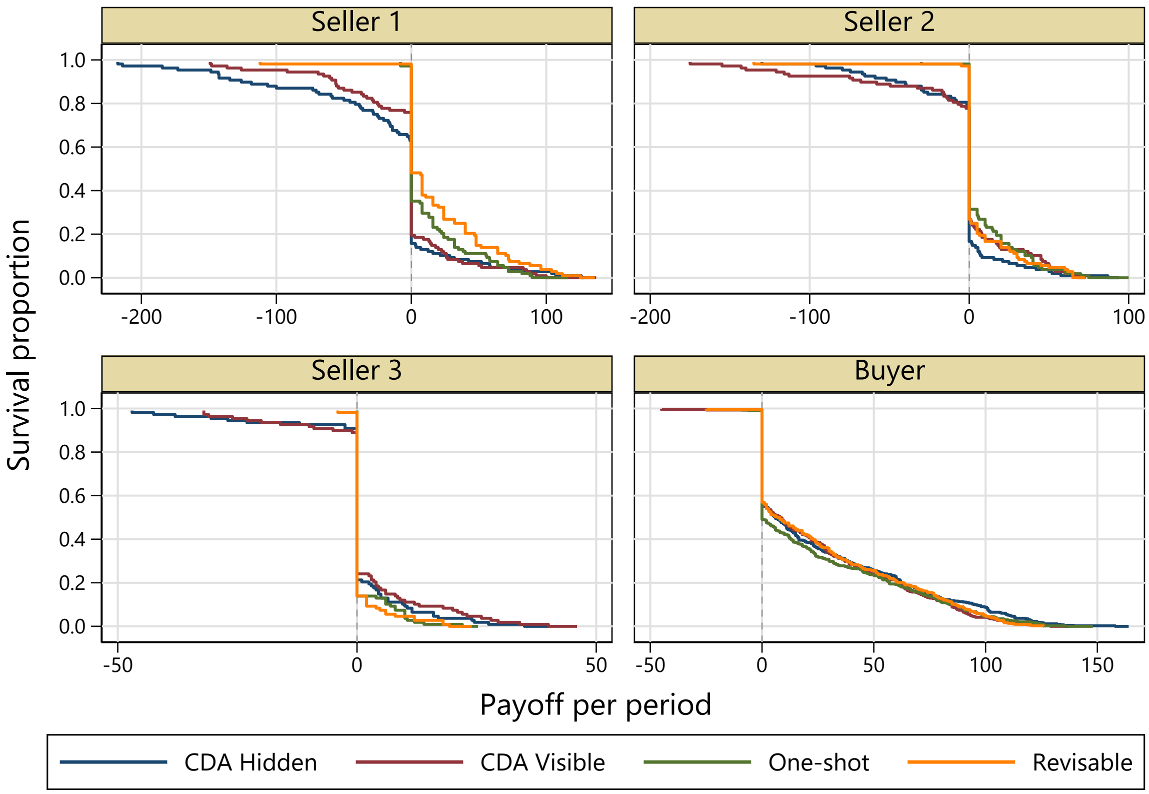

The figure plots empirical survival functions of payoffs by trader role and treatment. At each payoff level, a curve shows the proportion of observations with a higher payoff (a higher curve is better). Each seller observation is the payoff of the relevant seller role in a market-period; seller roles are identified by capacity because capacities rotate across participants. The buyer panel pools the individual payoffs of all four buyers, so each observation is a buyer-period. Notice the treatment effects were most pronounced for seller 1, with fewer losses and larger positive payoffs in the ONE-SHOT SCHEDULES and REVISABLE SCHEDULES treatments compared to the CDA treatments.