Structured-Bid Smart Auction: Exact Clearing with Exclusive-Choice and All-or-None Declarations

Ali Haydar Özer
Marmara University
Istanbul, Turkey
haydar.ozer@marmara.edu.tr

***Abstract*—Two-sided markets for heterogeneous indivisible items couple allocation decisions through seller reservation prices, buyer budgets, and buyer item caps. Smart Auction (SA) clears compact item-wise reports. We introduce Structured-Bid Smart Auction (SBSA), an exact pay-as-bid extension with exclusive-choice (EC) and all-or-none (AON) declarations. SBSA optimization jointly determines the allocation under these declarations and the ordinary market constraints. We evaluate it against structure-oblivious SA and optimal deletion-only repair of the SA allocation. The main study comprises 120 synthetic markets, 1,080 structured variants, and 2,280 allocation solves. When all eligible buyers received declarations, the mean per-market fraction of structured buyers with a structure-oblivious violation ranged from 1.3 to 19.7% across the representative treatments. SBSA retained 98.7–99.7% of the relaxation objective, and its advantage over repair was 0.6–4.2% of feasible SBSA gain from trade. In a separate experiment with up to 2,000 items, median wall times at the largest size ranged from 2.29 to 5.90 seconds across SA and SBSA.**

***Index Terms*—two-sided multi-item auctions, structured bids, exclusive-choice sets, all-or-none groups, budget constraints, pay-as-bid clearing, mixed-integer programming**

I. Introduction

Many two-sided markets allocate heterogeneous and indivisible items in one clearing round. Sellers may impose item-specific reservation prices, while buyers may have hard spending limits and limits on the number of acquired items. These restrictions couple decisions across items because accepting one bid can reduce the buyer's capacity to accept another. Related studies include unit-demand multi-item auctions [1] and bundle-based combinatorial allocation [2, 3], while hard budgets add further feasibility and mechanism-design difficulties [4–6]. In two-sided exchange, efficiency, incentives, and budget balance also cannot generally be treated as independent design objectives [7–9].

Smart Auction (SA) addresses the allocation coupling by clearing item-wise bids in one market-wide optimization problem with seller reserves, buyer budgets, buyer item caps, and single assignment of each item [10, 11]. The item-wise report is compact and treats every reserve-passing item as independently acceptable once the ordinary budget and item cap constraints are satisfied. General package bids can represent richer cross-item values, but they change the report language and lead to combinatorial winner determination [2, 3, 12].

This paper extends SA with two hard feasibility declarations. An *exclusive-choice* (EC) set permits zero or one selected item. An *all-or-none* (AON) set requires either every item or none. Structured-Bid Smart Auction (SBSA) maximizes reported gain from trade (GFT) while enforcing these declarations jointly with seller reserves, buyer budgets, item caps, and single assignment. We call this allocation problem *SBSA optimization*.

We evaluate SBSA optimization against two comparison baselines: (i) structure-oblivious SA and (ii) optimal deletion-only repair. Structure-oblivious SA ignores EC and AON declarations and can therefore return a structurally infeasible allocation. Its objective is an upper-bound benchmark. Optimal deletion-only repair selects the best structured-feasible subset of that SA allocation. The comparison measures how much feasible GFT SBSA preserves when it selects trades jointly under the declared structures rather than retaining only trades from the structure-oblivious allocation.

The evaluation addresses the following research questions:

RQ1: How often and how severely does structure-oblivious SA violate declared EC and AON requirements?

RQ2: How much feasible reported GFT does SBSA optimization retain relative to the structure-oblivious relaxation, and how does it compare with optimal deletion-only repair?

RQ3: How do representative EC-only, AON-only, and joint declarations behave as structure density and acquisition opportunity change, and how do the results change when group size and the buyer pool are matched?

RQ4: What model-size and solution-time overhead accompanies exact structured clearing over the evaluated market sizes?

The contributions are as follows:

(i) We define a structured extension of the SA report language and a compact exact SBSA optimization model that jointly enforces EC sets, overlapping AON sets, and their interactions with the ordinary market constraints.

(ii) We establish the objective ordering between SBSA optimization and two comparison baselines, i.e. structure-oblivious SA and optimal deletion-only repair, through their nested feasible regions.

(iii) We provide an evaluation with outcome-independent structure generation, a common eligible-buyer pool, nested densities, and base-market-level statistical inference.

(iv) We quantify structural violations, feasible-GFT retention

under SBSA optimization, its advantage over repair, and computational scaling, including an EC–AON sensitivity analysis with matched group sizes and buyer pools.

Within the evaluated settings, the mean per-market fraction of structured buyers with a structure-oblivious SA violation ranged from 1.3 to 19.7% at full eligible-pool density in the representative treatments. SBSA retained 98.7–99.7% of the structure-oblivious relaxation GFT and exceeded deletion-only repair by 0.6–4.2% of feasible SBSA GFT. In the sensitivity analysis with matched group sizes and buyer pools, the tested AON groups produced larger violation rates and larger SBSA advantages over repair than the tested EC groups at each size. The separate scalability grid reached 2,000 items, with median wall times of 2.29–5.90 seconds across SA and SBSA at the largest size.

The remainder of the paper positions SBSA within related work, recalls the required SA allocation model, presents the structured formulation, and reports the experimental evidence.

II. Related Work

SBSA lies between item-wise multi-item clearing and fully combinatorial bidding. In a combinatorial auction, bidders can attach values to bundles so that complements and substitutes are represented directly. Standard treatments cover bundle-based winner determination and its computational trade-offs [2, 3, 12]. OR/XOR and logical bidding languages make these relations more concise by specifying whether atomic bids may be accepted together, at most one may be accepted, or more general formulas must hold [13, 14]. Winner determination remains difficult in general, motivating restrictions on allowable bundles [15], linear-programming-based methods [2], and exact search algorithms [12]. Two-sided combinatorial exchanges extend bundle-based winner determination to bids and asks, multiple units, and exchange feasibility [9, 16].

An EC set has the same allocation-side effect as XOR over its member singleton bids, while an AON set resembles an optional AND/package relation. The analogy is limited. SBSA does not require a value for each package or an arbitrary nonadditive valuation. It retains one monetary bid per buyer–item edge, charges selected edges under pay-as-bid (PAB) clearing, permits items outside a declaration to be acquired independently, and allows multiple declarations for the same buyer. Overlapping AON sets are normalized into equality components rather than interpreted as separate package bids. SBSA is therefore less expressive than a general combinatorial bidding language, but it is compact and backward-compatible with the original SA report.

Package bidding is often motivated by the exposure problem. Item-wise clearing can leave a bidder with only part of a complementary collection, whereas package bids condition acceptance on the bundle [17]. SBSA addresses this concern through hard feasibility declarations. AON blocks partial acquisition of a declared group, and EC prevents the buyer from receiving multiple declared alternatives. These declarations combine hard allocation feasibility with additive item-wise bids.

Budget-constrained auction models cover homogeneous multi-unit goods [5] and heterogeneous items with additive valuations [6]. Polyhedral clinching provides incentive-compatible, individually rational, and Pareto-optimal mechanisms for divisible goods with single-parameter linear valuations, public budgets, and polymatroidal allocation constraints [18]. Hirai and Sato extend this approach to two-sided markets, assuming public buyer budgets and seller reserve prices equal to true valuations [19]. These guarantees rely on the respective model assumptions. AON feasibility is not downward-closed, that is, a complete group may be feasible although every nonempty proper subset is infeasible. The present paper focuses on exact allocation under PAB clearing. Relative to prior SA [10, 11], the contribution is the compact integration of restricted EC/AON feasibility and its evaluation against structure-oblivious SA and deletion-only repair.

III. Baseline Smart Auction

The full SA model and its broader mechanism set are given in the earlier Smart Auction studies [10, 11]. We recall only the exact pay-as-bid allocation model needed for SBSA.

A. Market, reports, and feasible allocation

Let $\mathcal{S}$ and $\mathcal{B}$ denote disjoint sets of sellers and buyers. Seller $s \in \mathcal{S}$ owns a set $\mathcal{I}_s$ of indivisible items. The item sets are disjoint, and $\mathcal{I} = \bigcup_{s\in\mathcal{S}} \mathcal{I}_s$. Item i has a seller-specified integer reservation price $r_i \in \mathbb{Z}_{\geq 0}$. Buyer $j \in \mathcal{B}$ reports an integer spending limit $s_j \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$, an item cap $k_j \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$, and integer item-wise bids b_{ji}. We write $\mathcal{E}_j \subseteq \mathcal{I}$ for the items on which the buyer submits a reserve-passing bid, so

$$b_{ji} \geq r_i, \qquad i \in \mathcal{E}_j. \tag{1}$$

Only submitted bids satisfying (1) are eligible for allocation.

For every $j \in \mathcal{B}$ and $i \in \mathcal{E}_j$, let $x_{ji} = 1$ when item i is assigned to buyer j, and $x_{ji} = 0$ otherwise. The baseline feasible set $\mathcal{F}_{\mathrm{SA}}$ is defined by

$$\sum_{j:i\in\mathcal{E}_j} x_{ji} \leq 1, \qquad i \in \mathcal{I}, \tag{2}$$

$$\sum_{i\in\mathcal{E}_j} x_{ji} \leq k_j, \qquad j \in \mathcal{B}, \tag{3}$$

$$\sum_{i\in\mathcal{E}_j} b_{ji}x_{ji} \leq s_j, \qquad j \in \mathcal{B}, \tag{4}$$

$$x_{ji} \in \{0,1\}, \qquad j \in \mathcal{B}, \quad i \in \mathcal{E}_j. \tag{5}$$

A constraint with an infinite right-hand side is omitted. Constraint (2) permits an item to remain unsold. Under PAB clearing, a selected buyer–item edge is charged b_{ji}. Therefore, (4) is the realized payment constraint, not an approximation.

B. Reported gain from trade objective

The allocation objective is reported gain from trade relative to seller reserves:

$$\phi(x) = \sum_{j\in\mathcal{B}} \sum_{i\in\mathcal{E}_j} (b_{ji} - r_i)x_{ji}. \tag{6}$$

This quantity is called reported generalized surplus in the preceding SA papers. Baseline SA selects an optimal allocation

$$x^{\mathrm{SA}} \in \arg\max_{x \in \mathcal{F}_{\mathrm{SA}}} \phi(x). \tag{7}$$

IV. Structured-Bid Smart Auction Model

A. Structured report language

SBSA attaches hard feasibility declarations to the buyer's eligible item-wise bids. Buyer j reports

$$\rho_j = \left(\{b_{ji}\}_{i \in \mathcal{E}_j},\ s_j,\ k_j,\ \mathcal{C}_j,\ \mathcal{A}_j\right), \tag{8}$$

where $\mathcal{C}_j$ is a finite family of nontrivial EC sets and $\mathcal{A}_j$ is a finite family of nontrivial AON sets, each contained in $\mathcal{E}_j$. For an allocation x, buyer j receives

$$T_j(x) = \{i \in \mathcal{E}_j : x_{ji} = 1\}. \tag{9}$$

The declarations restrict which subsets of the reported item-wise bids may be accepted. They do not change the bid or reserve associated with an accepted edge.

B. Small example

A short example helps distinguish the two structures and illustrates why they are naturally modeled as hard feasibility declarations rather than as post-processing rules. Consider one buyer with item cap $k_1 = 3$ and spending limit $s_1 = 170$. She submits bids on four items. The seller reserve prices and the buyer bids are shown in Fig. 1. The buyer declares the EC set $C_1 = \{a, b\}$ and the AON set $G = \{c, d\}$.

The complete reported feasible family is

$$\{\emptyset, \{a\}, \{b\}, \{c, d\}, \{a, c, d\}, \{b, c, d\}\}.$$

The bundle $\{a, b\}$ violates the EC condition. Bundles $\{c\}$, $\{d\}$, and $\{a, c\}$ are examples of partial AON bundles and are therefore infeasible. If bids equal true item values, the additive value of $\{c, d\}$ is $50 + 45 = 95$.

C. Exclusive-choice declarations

For each $C \in \mathcal{C}_j$, EC semantics require

$$|T_j(x) \cap C| \le 1. \tag{10}$$

Thus, the buyer can receive no member or one member of C, but not two or more. The corresponding mixed-integer programming (MIP) constraint is

$$\sum_{i \in C} x_{ji} \le 1, \qquad j \in \mathcal{B},\ C \in \mathcal{C}_j. \tag{11}$$

Overlapping EC sets are permitted, and every declared inequality is enforced.

D. All-or-none declarations and normalization

For each declared $G \in \mathcal{A}_j$, AON semantics require

$$T_j(x) \cap G \in \{\emptyset, G\}. \tag{12}$$

A buyer may receive items outside G, subject to the remaining constraints. When AON sets overlap, enforcing each set separately implies transitive equality over their union. SBSA represents each connected overlap component directly.

Let $U_j = \bigcup_{G \in \mathcal{A}_j} G$. For $i, i' \in U_j$, define $i \sim_j i'$ when there are $G_1, \dots, G_q \in \mathcal{A}_j$ satisfying

$$\begin{aligned} i \in G_1, \qquad & i' \in G_q, \\ G_\ell \cap G_{\ell+1} \neq \emptyset, \qquad & \ell = 1, \dots, q-1. \end{aligned} \tag{13}$$

The equivalence classes form the normalized AON components, denoted by $\mathcal{H}_j$. For each $H \in \mathcal{H}_j$, choose an anchor $i_0(H)$ and add

$$x_{ji} = x_{ji_0(H)}, \qquad i \in H \setminus \{i_0(H)\}. \tag{14}$$

These equalities are equivalent to the raw AON declarations. A component is either entirely selected or entirely unselected. A component of size $|H|$ needs $|H| - 1$ equalities and no additional AON binary variable. The choice of anchor does not affect the feasible allocations.

E. EC–AON interactions

EC and AON declarations are enforced jointly. If an EC set contains at least two items from one normalized AON component, selecting that component would select at least two members of the EC set. The component is therefore forced off:

$$|C \cap H| \ge 2 \quad \Longrightarrow \quad x_{ji_0(H)} = 0, \qquad C \in \mathcal{C}_j,\ H \in \mathcal{H}_j. \tag{15}$$

This is a logical consequence of the combined declarations, not a rejection of the report. When an EC set intersects a component in at most one item, Constraints (11) and (14) apply directly. In particular, one EC set may couple several distinct AON components and permit at most one of the represented alternatives.

F. SBSA optimization

SBSA optimization determines an allocation directly from the submitted item-wise bids, seller reserves, buyer budgets and item caps, and EC/AON declarations. Let $\mathcal{F}_{\mathrm{S}}$ be the subset of $\mathcal{F}_{\mathrm{SA}}$ satisfying (11), (14), and the implied fixes in (15):

$$\mathcal{F}_{\mathrm{S}} = \{x \in \mathcal{F}_{\mathrm{SA}} : x \text{ satisfies all declared structures}\}. \tag{16}$$

The SBSA allocation and its objective value are

$$\begin{aligned} x^{\mathrm{S}} &\in \arg\max_{x \in \mathcal{F}_{\mathrm{S}}} \phi(x), \\ \Phi_{\mathrm{S}} &= \phi(x^{\mathrm{S}}). \end{aligned} \tag{17}$$

All eligible buyer–item allocation variables are decided jointly, with the ordinary market constraints and the declared structures enforced in the same MIP. If all EC and AON families are empty, $\mathcal{F}_{\mathrm{S}} = \mathcal{F}_{\mathrm{SA}}$ and SBSA optimization reduces to the baseline SA allocation problem.

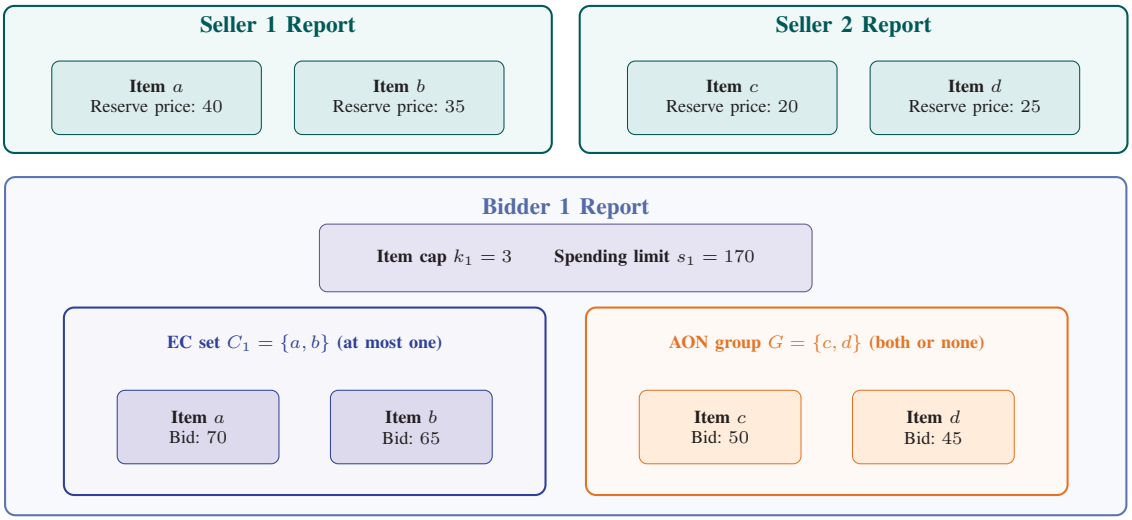


Fig. 1. Small example of a structured report. Seller 1 offers items a and b, while Seller 2 offers items c and d. Each seller declares a reserve price for each item. Bidder 1 has item cap $k_1 = 3$ and spending limit $s_1 = 170$. The bidder declares $C_1 = \{a, b\}$ as an exclusive-choice (EC) set, so at most one of these items can be allocated to her. She also declares $G = \{c, d\}$ as an all-or-none (AON) set, so either both items are allocated or neither is allocated.

G. Complexity and formulation size

Exact SBSA optimization is NP-hard even when $\mathcal{C}_j = \mathcal{A}_j = \emptyset$. With one buyer, zero reserves, no item cap, and integer bids, maximizing $\sum_i b_i x_i$ subject to $\sum_i b_i x_i \le s$ is the optimization form of Subset Sum. Deciding whether the optimum equals s decides a Subset Sum instance [20, 21].

The structured extension remains compact. If N_{off} is the number of forced-off AON components, the number of structured rows is

$$m_{\text{struct}} = \sum_{j \in \mathcal{B}} |\mathcal{C}_j| + \sum_{j \in \mathcal{B}} \sum_{H \in \mathcal{H}_j} (|H| - 1) + N_{\text{off}}. \tag{18}$$

Thus, SBSA adds one row per EC set, $|H|-1$ equality rows per normalized AON component, and one fixing row per forced-off component, but no structured binary variables.

H. Comparison baselines and objective ordering

We compare SBSA optimization with two baselines: (i) structure-oblivious SA and (ii) optimal deletion-only repair.

a) Structure-oblivious SA: SA ignores all EC and AON declarations and solves the baseline model once per base market. The same optimal SA allocation is evaluated against every structured variant derived from that market. Its objective value is

$$\Phi_{\text{SA}} = \phi(x^{\text{SA}}). \tag{19}$$

If $x^{\text{SA}} \notin \mathcal{F}_{\text{S}}$, Φ_{SA} is the *structure-oblivious relaxation objective*. It is attained by an infeasible allocation and serves only as an upper-bound benchmark for the declared structured problem.

b) Optimal deletion-only repair: Let $\widehat{x}^{\text{SA}} = x^{\text{SA}}$ denote this returned SA allocation. The repair baseline selects its best structured-feasible subset by optimizing over

$$\mathcal{F}_{\text{R}}(\widehat{x}^{\text{SA}}) = \{x \in \mathcal{F}_{\text{S}} : x_{ji} \le \widehat{x}^{\text{SA}}_{ji} \ \forall j \in \mathcal{B},\ i \in \mathcal{E}_j\}. \tag{20}$$

Its solution and objective are

$$\begin{aligned} x^{\text{R}} &\in \operatorname*{arg\,max}_{x \in \mathcal{F}_{\text{R}}(\widehat{x}^{\text{SA}})} \phi(x), \\ \Phi_{\text{R}} &= \phi(x^{\text{R}}). \end{aligned} \tag{21}$$

Repair can retain or delete trades from the reference SA allocation. Its optimum is therefore conditional on the particular optimal SA allocation returned.

Proposition 1. *If SBSA optimization and both baseline allocation MIPs are solved to optimality, then*

$$\Phi_{\text{R}} \le \Phi_{\text{S}} \le \Phi_{\text{SA}}. \tag{22}$$

Proof: By construction, $\mathcal{F}_{\text{R}}(\widehat{x}^{\text{SA}}) \subseteq \mathcal{F}_{\text{S}} \subseteq \mathcal{F}_{\text{SA}}$. Maximizing the same objective ϕ over nested feasible sets gives the two inequalities. The all-zero allocation belongs to every set, so the repair problem is feasible. ■

For true item values v_{ji}, a buyer's value on bundles satisfying her true declarations is $V_j(T) = \sum_{i \in T} v_{ji}$, with utility $V_j(T) - P_j$ at total payment P_j, subject to her true budget and cap. Violating bundles are inadmissible. Taking reserves as seller opportunity costs, true structured welfare gain is $W(x) = \sum_{j,i}(v_{ji} - r_i)x_{ji}$ on feasible allocations only. Under truthful reports, $W(x) = \phi(x)$ on this domain. Thus, $\Phi_{\text{SA}} - \Phi_{\text{S}}$ is a reduction from a relaxation, not a welfare loss when SA is inadmissible. In contrast, $\Phi_{\text{S}} - \Phi_{\text{R}}$ measures the feasible reported-GFT advantage over deletion-only repair of the reference allocation.

V. Experimental Evaluation

A. Questions, methods, and experimental design

The evaluation follows RQ1–RQ4 from Section I. The paired main and sensitivity studies evaluate SBSA optimization, defined in Section IV-F, against the structure-oblivious SA and optimal deletion-only repair baselines from Section IV-H.

Each structured variant is solved by SBSA using its market reports and declarations. The corresponding repair benchmark is computed from the base market's SA allocation. A separate scalability grid compares SA and SBSA only. We use truthful PAB reports as an allocation-benchmark assumption to isolate allocation and feasibility effects.

Table I summarizes the main experimental grid. All experiments were conducted on a system with an Apple M1 Pro chip (10-core CPU) and 16 GB RAM, running macOS version 26.5. Gurobi Optimizer (version 13.0) was used as the MIP solver. The grid contains three nonregular profiles (*clustered*, *stress*, and *budget-lite*) with multiple categories and heterogeneous budgets. Every base market has 200 items and either 100 or 200 buyers, giving buyer-to-item ratios 0.5 and 1.0. The target mean numbers of interested items per buyer are 32, 48, and 32 for the three profiles, respectively. We combine $\beta_{\mathrm{budget}} \in \{0.50, 0.80\}$ with both ratios, use item caps 5–8, and evaluate ten base markets per profile–budget–ratio combination. Here, β_{budget} is the mean of the per-buyer beta draw used to construct spending limits, not a monetary budget average. The main grid contains 120 base markets and 1,080 structured variants, requiring 2,280 allocation solves.

All allocation solves use one thread, a 600 s limit, and a primary-gap target of zero with tolerance 10^{-12}. The objective hierarchy maximizes GFT, then accepted bid volume, then trade count. The returned SA allocation is reused for all variants of a market.

Structured declarations are generated before any allocation is solved. For buyer j, candidate items come from $E_j^+ = \{i \in \mathcal{E}_j : b_{ji} > r_i\}$. A preferred subset of candidate items is searched first, and the search is expanded to all of E_j^+ when necessary. Primary EC sets contain four same-category items and permit at least two alternatives under the ordinary budget and cap. Primary AON sets contain two jointly affordable cross-category items and prefer different sellers. A joint declaration uses disjoint EC and AON sets and permits the complete AON pair together with one EC alternative. Categories are synthetic proxies for substitutes and complements.

The common eligible set consists of buyers for whom a valid joint candidate exists. One common ordering of this set is used for every mode, and the density selections are nested. The 0.25 selection is contained in the 0.50 selection, which is contained in the 1.00 selection. Treatment density is defined relative to this eligible set. The realized structured-buyer share among all buyers is reported separately. Hence, the primary EC-only, AON-only, and joint variants use the same selected buyers and reuse the corresponding declaration.

The matched-cardinality sensitivity analysis uses the *stress* profile with 250 items and 250 buyers, $\beta_{\mathrm{budget}} = 0.50$, density 0.50, sizes two to four, ten base markets, and the intersection of buyers eligible for nested EC and AON groups. The separate scalability grid uses the *stress* profile with 100–2,000 items, buyer-to-item ratios 1 and 2, $\beta_{\mathrm{budget}} = 0.50$, target mean interest 32, item caps 5–8, and joint EC+AON declarations at densities 0.50 and 1.00. Five base markets are evaluated at each market-size and buyer-to-item-ratio combination, yielding 150 SA and SBSA solves.

B. Metrics and statistical analysis

Among the three compared methods, EC and AON violations can occur only in structure-oblivious SA. Optimal deletion-only repair and SBSA enforce all declarations. The following structural metrics are evaluated on the SA allocation. For each EC set C, we record its selected-member count, activation ($\sum_{i \in C} x_{ji} \geq 1$), violation ($\sum_{i \in C} x_{ji} > 1$), and excess selected members. For each normalized AON component H, a partial selection, $0 < \sum_{i \in H} x_{ji} < |H|$, is a violation. The primary structural metric is the fraction of structured buyers with at least one violated declaration. EC activation and two buyer-level diagnostics provide context. Multi-win exposure is the fraction of EC-structured buyers receiving at least two SA items in total, and conditional violation is the fraction of exposed buyers who violate their EC set. These diagnostics do not replace the unconditional violation rate, and multi-win exposure does not assert that the two items belong to the EC set. Conditional violation is recorded as missing when its denominator is zero.

The economic ratios are

$$\begin{aligned} R_{\mathrm{S}} &= \frac{\Phi_{\mathrm{S}}}{\Phi_{\mathrm{SA}}}, \qquad R_{\mathrm{R}} = \frac{\Phi_{\mathrm{R}}}{\Phi_{\mathrm{SA}}}, \\ E_{\mathrm{R}} &= \frac{\Phi_{\mathrm{R}}}{\Phi_{\mathrm{S}}}, \qquad G_{\mathrm{adv}} = \frac{\Phi_{\mathrm{S}} - \Phi_{\mathrm{R}}}{\Phi_{\mathrm{S}}}. \end{aligned} \tag{23}$$

R_{S} and R_{R} measure GFT retention relative to the SA relaxation, while E_{R} is repair efficiency relative to SBSA. We also record the absolute SBSA advantage over repair, $\Delta_{\mathrm{adv}} = \Phi_{\mathrm{S}} - \Phi_{\mathrm{R}}$. Its normalized counterpart, G_{adv}, expresses this difference as a fraction of feasible SBSA GFT, not as a percentage increase relative to repair GFT. Ratios with a zero denominator are stored as missing. When SA violates the declarations, Φ_{SA} and ratios using it refer to the structure-oblivious relaxation, not to feasible structured welfare.

All planned markets and variants were included, and every allocation MIP solve reached proven optimality. No result was excluded from the structural, economic, or computational analyses.

The experimental unit is a base market, defined as the generated market before EC/AON declarations are added. For each treatment, structural percentages and economic ratios are computed separately for each market and then averaged across markets. The selected-member histogram in Section V-C instead pools EC sets. We report means with 95% confidence intervals (CIs) obtained from 10,000 stratified cluster-bootstrap resamples. Whole markets are resampled with replacement within groups defined by profile, budget level, and buyer-to-item ratio, keeping all methods and variants of a market together. The same resamples are reused across treatments. Runtime summaries report medians, interquartile ranges (IQRs), 90th percentiles, and maxima.

C. RQ1: Structural violations under structure-oblivious clearing

Fig. 2 shows the unconditional fraction of structured buyers with any SA violation and the EC diagnostics. SBSA optimiza-

TABLE I
EXPERIMENTAL CONFIGURATION

Market size	200 items with 100 or 200 buyers (ratios 0.5, 1.0)	**Mean interests**	32 (clustered/budget-lite), 48 (stress)
Budget/cap settings	$\beta_{\text{budget}} = 0.50, 0.80$, item caps 5–8	**Value generation**	reserve scale 1.0, margin scale 2.5
Structured treatments	EC-only, AON-only, joint at densities 0.25, 0.50, 1.00	**Primary declarations**	one EC set of size 4 and/or one AON set of size 2
Main-grid counts	2,280 allocation solves	**Clearing**	PAB with zero auctioneer share and reported-GFT maximization

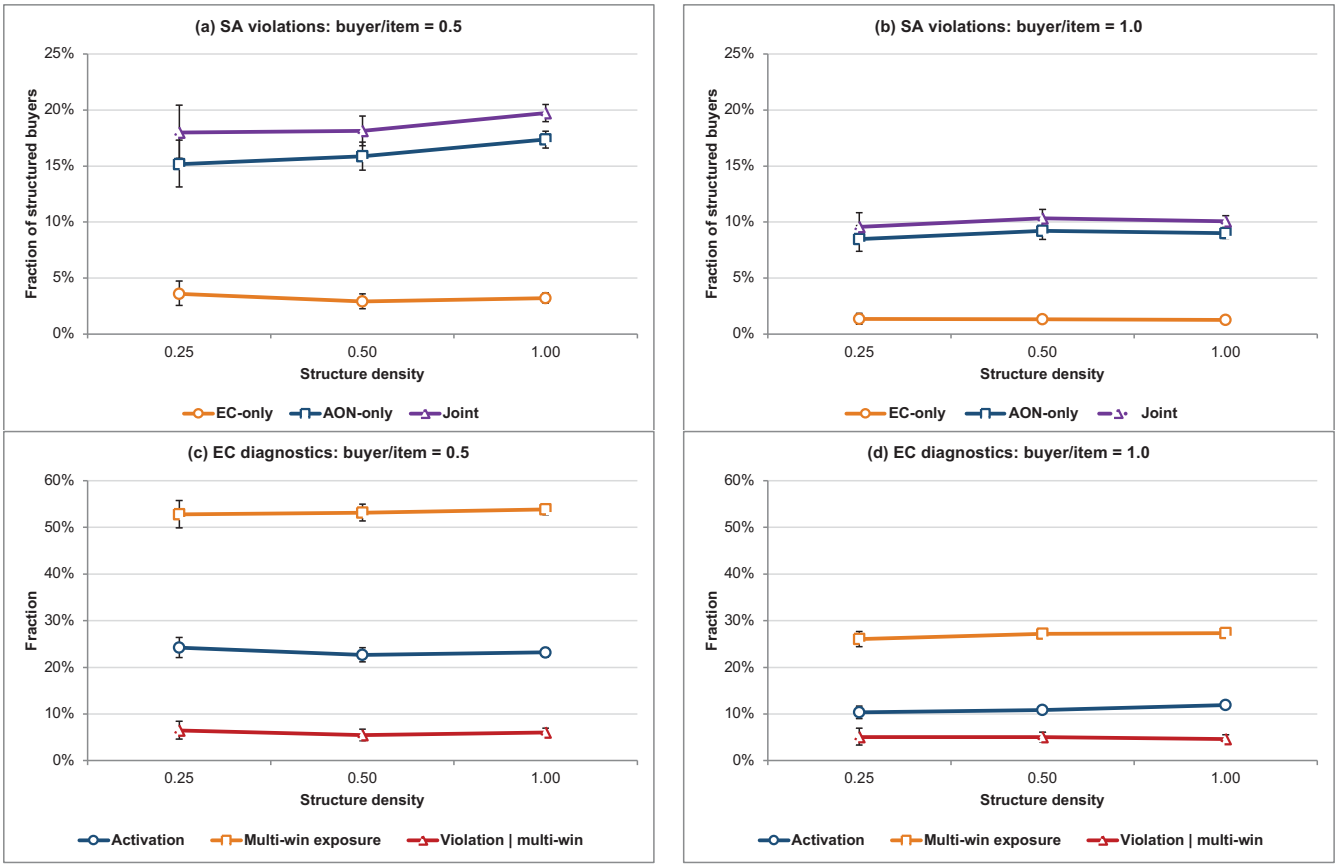


Fig. 2. Structure-oblivious SA outcomes. Panels (a)–(b) report the mean per-market fraction of structured buyers with at least one violated declaration. Panels (c)–(d) report mean per-market EC activation, multi-win exposure (at least two total SA items), and EC violation conditional on multi-win exposure. EC declarations and the fixed SA allocation are common to EC-only and joint variants at a given density, so these diagnostics are identical and are plotted once. Error bars are 95% stratified cluster-bootstrap CIs (60 base markets at each buyer/item ratio).

tion and the repair baseline both satisfy all declared structures by construction. In the more item-abundant ratio-0.5 markets, the mean per-market violation fractions across densities were 2.9–3.6% for EC-only, 15.2–17.4% for AON-only, and 18.0–19.7% for joint declarations. At ratio 1.0, the corresponding ranges were 1.3–1.4%, 8.5–9.2%, and 9.6–10.3%. The EC-only result is weak and not monotone in density. Its confidence intervals overlap strongly. AON-only and joint violations are more frequent in these representative treatments, especially at ratio 0.5, but this descriptive difference also includes the EC=4 versus AON=2 cardinality choice and the family-specific group-construction rules.

The diagnostics explain part of the small EC incidence. These percentages are means of per-market fractions. At ratio 0.5, EC activation was 22.7–24.2%, although multi-win exposure was 52.8–53.9%. Among exposed buyers, conditional violation was 5.5–6.5%. At ratio 1.0, activation was 10.4–11.9%, multi-win exposure was 26.0–27.3%, and conditional violation was 4.6–5.1%. Thus, ordinary acquisition opportunity is necessary but not sufficient for an EC violation. The structure-oblivious allocation must also choose multiple members of the particular declared substitute set. In the pooled full-density selected-member histogram, 87.9% of violating ratio-0.5 EC sets had exactly two selected members and 12.1% had three. The corresponding ratio-1.0 fractions were 92.9% and 7.1%, and no violating EC set had all four members selected. Because the primary AON declarations are pairs, every partial AON violation contains exactly one selected and one missing

member.

D. RQ2: Feasible GFT under SBSA optimization

Table II reports the density-1.00 outcomes, where all buyers in the common eligible pool receive declarations. Across the six representative rows, mean SBSA retention is 98.7–99.7% of the structure-oblivious relaxation objective. Repair efficiency relative to SBSA is 95.8–99.4%, so the normalized SBSA advantage over repair is 0.6–4.2%. SBSA strictly improves on repair in 83.3–100% of the variants. The EC-only advantage is small but not uniformly zero. AON-only and joint treatments show a larger SBSA advantage over repair in this design.

Fig. 3 gives all density levels. Mean SBSA retention over the 18 displayed combinations is 98.73–99.98%, whereas deletion-only repair retains 94.60–99.91% of the relaxation objective. The SBSA advantage over repair grows most clearly with density for the representative AON-only and joint treatments. This gap measures the additional feasible reported GFT obtained by jointly selecting trades under the declared structures rather than selecting only a subset of the structure-oblivious SA allocation.

E. RQ3: Density, acquisition opportunity, and matched cardinality

The primary grid shows two main patterns. First, the ratio-0.5 markets create more opportunities to acquire multiple items and have larger structural violation rates than the ratio-1.0 markets. Second, EC violation rates do not change monotonically with density, although the normalized SBSA advantage over repair for EC increases from 0.14% to 0.67% at ratio 0.5 and from 0.07% to 0.58% at ratio 1.0. For the representative AON-only and joint treatments, higher density generally increases the SBSA advantage over repair, but the ratio-1.0 AON violation rate is nearly flat from density 0.50 to 1.00. Thus, the unconditional EC result is weak and nonmonotone, while the conditional EC statistics remain diagnostic only.

Table III reports a sensitivity analysis with matched group sizes and buyer pools. At each group size, the paired treatments use the same selected buyers, but the item groups and validity conditions remain family-specific (same-category EC groups versus cross-category AON groups). Under these representative constructions, AON produces a larger violation rate and a larger SBSA advantage over repair than EC at sizes two, three, and four. Both effects increase with AON cardinality, whereas the EC effects remain small. These are treatment-specific findings from ten paired markets at each size, not a universal ranking of EC and AON semantics. In the family-specific nested sensitivity analysis, mean EC violation rates range from 0.32 to 1.28% and normalized SBSA advantages over repair from 0.01 to 0.16% across sizes two to four, with 50% of all buyers selected. The AON family-specific and matched-nested violation and GFT outcomes coincide for these markets.

F. RQ4: Computational behavior and scalability

The compact formulation adds constraints but no structured binary variables. SA and SBSA had identical binary-variable counts for every main-grid market. Across the 60 base markets at each buyer-to-item ratio, the median counts were 3,213.5 at ratio 0.5 and 6,477.5 at ratio 1.0. For joint density 1.00, median constraint counts increased from 400 to 549 and from 600 to 884, respectively. Median SA wall times were 0.120 s (IQR 0.094–0.157) and 0.149 s (0.130–0.186). The corresponding joint SBSA medians were 0.122 s (0.103–0.163) and 0.182 s (0.157–0.226). For comparison, the optimal deletion-only repair baseline was faster in these rows, with medians 0.022 s and 0.039 s, consistent with its restricted retain-or-delete feasible region. Across all main-grid solves of each method, the 90th-percentile/maximum wall times were 0.242/1.001 s for SA, 0.050/0.054 s for repair, and 0.245/1.627 s for SBSA.

Fig. 4 reports the separate scalability experiment. At 2,000 items, median wall times for ratio 1 were 2.29 s for SA, 3.24 s for SBSA at density 0.50, and 5.90 s at density 1.00. For ratio 2, the corresponding medians were 3.69, 4.45, and 5.30 s. All 150 scalability MIP solves reached proven optimality. Across the complete scalability grid, the 90th-percentile wall time was 4.15 s, the maximum was 40.45 s, the median node count was 3, and the maximum was 274.

VI. Discussion

The structural results answer the central feasibility question within the evaluated settings. Ordinary budgets, item caps, and market competition did not eliminate EC and AON violations under structure-oblivious clearing. A structure-oblivious optimum can select several alternatives from one EC set or only part of an AON component. Such an allocation is optimal only for the relaxed item-wise report. SBSA optimization ensures structural feasibility by enforcing the declarations jointly with the ordinary market constraints during winner determination. It therefore treats structured feasibility as part of allocation, not as a post-clearing correction.

Optimal deletion-only repair is a demanding post-processing baseline. It chooses the best structured-feasible subset of the particular optimal SA allocation used as its reference. SBSA optimization instead chooses jointly among all eligible bids under the declared structures. Therefore, $\Phi_{\mathrm{S}} - \Phi_{\mathrm{R}}$ measures SBSA's feasible-GFT advantage over the repair baseline. That advantage is small in many EC cases but remains positive in most variants, and it is larger in the representative AON-only and joint cases. In contrast, $\Phi_{\mathrm{SA}} - \Phi_{\mathrm{S}}$ is only a reduction from a relaxation and is not a loss in true structured welfare.

The EC diagnostics clarify why EC effects are opportunity-dependent. Winning several items is necessary for an EC violation but is not sufficient. The allocation must activate the declared set and select more than one of its particular members. The ratio-1.0 markets offer fewer multi-win opportunities and show lower EC incidence. Even at ratio 0.5, many exposed buyers do not activate their EC set. Thus, these low rates do not establish that EC is unimportant.

The computational evidence is consistent with the compact formulation. Structured rows increase the constraint count without adding structured binary variables, and median overhead is modest in the 200-item main grid. Runtime grows with

TABLE II. DENSITY-1.00 OUTCOMES FOR SBSA AND ITS BASELINES (PERCENT, WITH 95% CIS WHERE APPLICABLE)

Buyer/item	Mode	Structured share	SA violation	R_{S}	E_{R}	G_{adv}	$\Phi_{\mathrm{S}} > \Phi_{\mathrm{R}}$
0.5	EC-only	75.1	3.2 [2.7, 3.7]	99.7 [99.6, 99.8]	99.3 [99.0, 99.6]	0.7 [0.4, 1.0]	88.3
0.5	AON-only	75.1	17.4 [16.6, 18.1]	99.0 [98.8, 99.2]	96.5 [95.8, 97.1]	3.5 [2.9, 4.2]	100.0
0.5	Joint	75.1	19.7 [19.0, 20.5]	98.7 [98.5, 98.9]	95.8 [95.0, 96.5]	4.2 [3.5, 5.0]	100.0
1.0	EC-only	74.0	1.3 [1.0, 1.5]	99.6 [99.3, 99.9]	99.4 [99.1, 99.7]	0.6 [0.3, 0.9]	83.3
1.0	AON-only	74.0	9.0 [8.5, 9.4]	99.1 [99.0, 99.2]	97.2 [96.9, 97.6]	2.8 [2.4, 3.1]	100.0
1.0	Joint	74.0	10.1 [9.6, 10.6]	98.7 [98.4, 99.0]	96.6 [96.2, 97.1]	3.4 [2.9, 3.8]	100.0

Entries are means of per-market values. Buyers are not pooled across markets, and each row summarizes 60 base markets. Structured share is the mean fraction of all buyers receiving declarations. Brackets give 95% CIs for the displayed means. $R_{\mathrm{S}} = \Phi_{\mathrm{S}}/\Phi_{\mathrm{SA}}$, $E_{\mathrm{R}} = \Phi_{\mathrm{R}}/\Phi_{\mathrm{S}}$, and $G_{\mathrm{adv}} = (\Phi_{\mathrm{S}} - \Phi_{\mathrm{R}})/\Phi_{\mathrm{S}}$. The final column is the percentage of variants in which SBSA strictly outperforms repair.

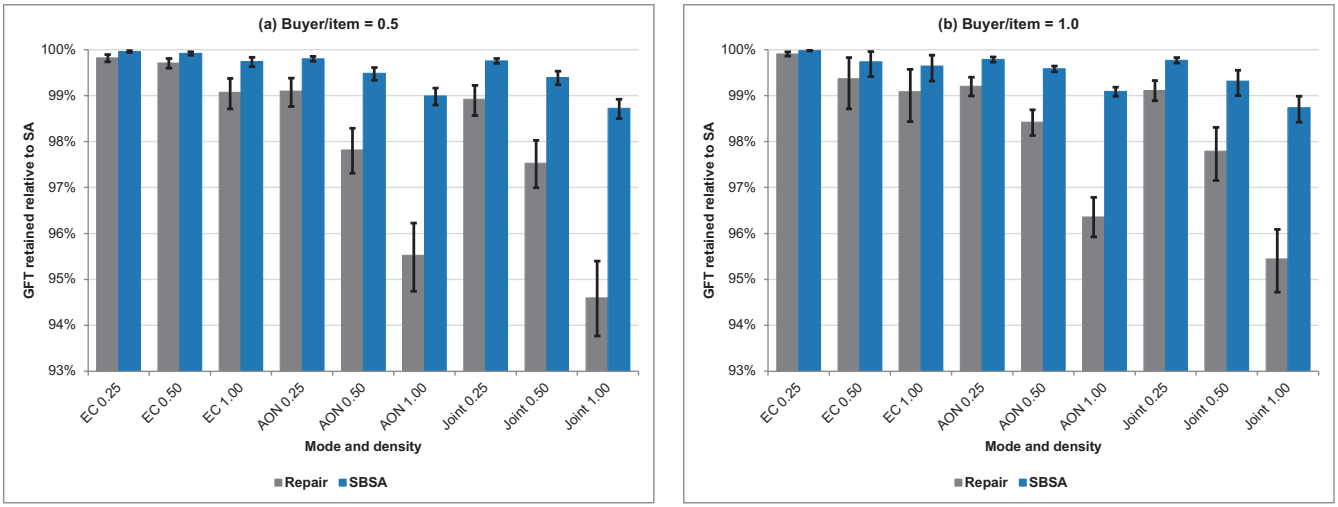


Fig. 3. Feasible GFT retained by SBSA optimization and the optimal deletion-only repair baseline, each divided by Φ_{SA}. When SA violates a declaration, the denominator is the structure-oblivious relaxation objective. Bars are means. Error bars are 95% stratified cluster-bootstrap CIs (60 base markets at each buyer/item ratio). The bar gap is the mean of $(\Phi_{\mathrm{S}} - \Phi_{\mathrm{R}})/\Phi_{\mathrm{SA}}$ across markets. The vertical axis is truncated to 93–100%.

TABLE III. CARDINALITY- AND BUYER-POOL-MATCHED SENSITIVITY (PERCENT, MEANS WITH 95% CIS)

Family	Size	SA violation	R_{S}	E_{R}	G_{adv}	$\Phi_{\mathrm{S}} > \Phi_{\mathrm{R}}$
EC	2	0.42 [0.00, 0.84]	100.00 [99.99, 100.00]	99.99 [99.98, 100.00]	0.01 [0.00, 0.02]	30
AON	2	7.46 [6.37, 8.50]	99.72 [99.60, 99.83]	99.23 [98.56, 99.69]	0.77 [0.31, 1.44]	100
EC	3	1.08 [0.39, 2.00]	99.96 [99.90, 100.00]	99.96 [99.90, 99.99]	0.04 [0.01, 0.10]	50
AON	3	10.49 [9.07, 12.14]	99.34 [99.13, 99.55]	98.41 [97.51, 99.26]	1.59 [0.74, 2.49]	100
EC	4	1.36 [0.52, 2.35]	99.88 [99.69, 99.99]	99.91 [99.81, 99.98]	0.09 [0.02, 0.19]	50
AON	4	12.80 [10.86, 14.84]	98.95 [98.61, 99.29]	98.05 [96.96, 99.04]	1.95 [0.96, 3.04]	100

The design uses the *stress* profile with 250 items and 250 buyers, $\beta_{\mathrm{budget}} = 0.50$, density 0.50, and a matched-nested eligible-buyer pool. Each row summarizes ten paired base markets. The matched eligible pool contains a mean 59.12% of all buyers (range 53.2–63.6%). The mean selected share is 29.44% after density rounding. Item groups and semantic validity conditions remain family-specific. G_{adv} is the normalized SBSA advantage over repair. The final column is the percentage of variants in which SBSA strictly outperforms repair.

market size and is higher at some dense SBSA points. The 2,000-item results demonstrate feasibility for the tested exact PAB instances.

VII. CONCLUSION

SBSA extends item-wise Smart Auction reports with hard EC and AON declarations. SBSA optimization jointly selects trades under these declarations, seller reserves, buyer budgets, item caps, and single assignment during exact PAB clearing. It determines the structured-feasible allocation directly, without requiring an SA allocation or a deletion-repair stage. In the evaluated settings, SBSA retains a high fraction of the structure-oblivious relaxation objective and preserves more feasible reported GFT than optimal deletion-only repair in most variants. EC effects depend on activation and multi-item acquisition opportunity, and the sensitivity analysis compares EC and AON using matched group sizes and buyer pools with family-specific construction rules. The compact formulation leaves the allocation-binary count unchanged and has modest median overhead in the 200-item grid. All MIP solves in the evaluation reached proven optimality, including the scalability cases with

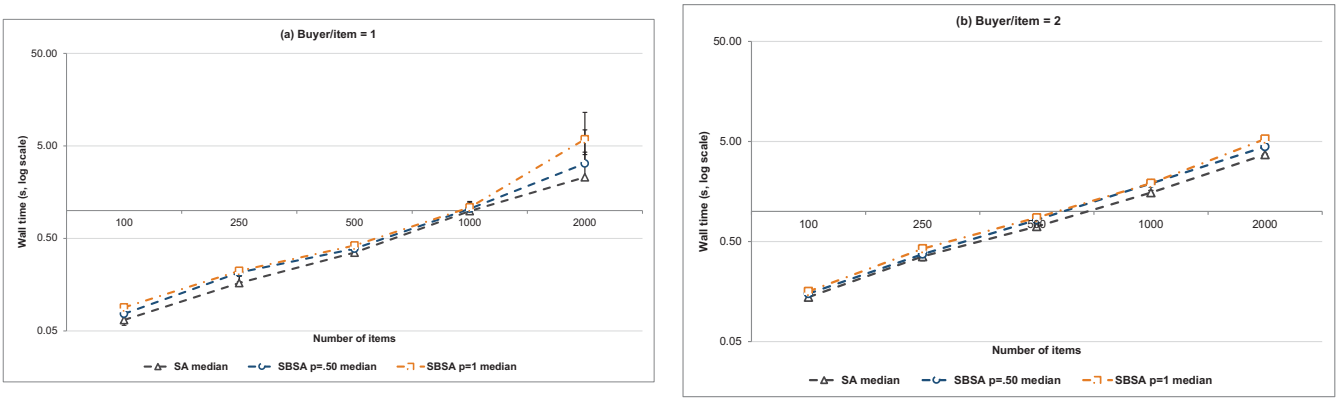


Fig. 4. Scalability of PAB clearing with joint EC=4 and AON=2. Panels use buyer-to-item ratios 1 and 2. SBSA densities are 0.50 and 1.00. Points are median wall times over five base markets, with 25th–75th percentile bars on a logarithmic vertical axis.

up to 2,000 items. These findings characterize the synthetic settings, report semantics, and market profiles studied here.

REFERENCES

[1] G. Demange, D. Gale, and M. Sotomayor, "Multi-item auctions," *Journal of Political Economy*, vol. 94, no. 4, pp. 863–872, 1986.

[2] N. Nisan, "Bidding and allocation in combinatorial auctions," in *Proceedings of the 2nd ACM Conference on Electronic Commerce*, ser. EC '00. New York, NY, USA: Association for Computing Machinery, 2000, pp. 1–12.

[3] P. Cramton, Y. Shoham, and R. Steinberg, Eds., *Combinatorial Auctions*. Cambridge, MA, USA: MIT Press, 2006.

[4] Y.-K. Che and I. Gale, "Standard auctions with financially constrained bidders," *The Review of Economic Studies*, vol. 65, no. 1, pp. 1–21, Jan. 1998.

[5] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi, "Multi-unit auctions with budget-constrained bidders," in *Proceedings of the 6th ACM Conference on Electronic Commerce*, ser. EC '05. New York, NY, USA: Association for Computing Machinery, 2005, pp. 44–51.

[6] P. Dütting, M. Henzinger, and M. Starnberger, "Auctions for heterogeneous items and budget limits," *ACM Transactions on Economics and Computation*, vol. 4, no. 1, pp. 1–17, 2015.

[7] R. B. Myerson and M. A. Satterthwaite, "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, vol. 29, no. 2, pp. 265–281, 1983.

[8] R. P. McAfee, "A dominant strategy double auction," *Journal of Economic Theory*, vol. 56, no. 2, pp. 434–450, 1992.

[9] D. C. Parkes, J. R. Kalagnanam, and M. Eso, "Achieving budget-balance with Vickrey-based payment schemes in exchanges," in *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*. San Francisco, CA, USA: Morgan Kaufmann, 2001, pp. 1161–1168.

[10] A. H. Özer and J. M. Ennis, "Smart Auction: Optimization-based two-sided multi-item auctions with budgets, item caps, and reservation prices," *Electronic Commerce Research and Applications*, 2026, under review.

[11] ——, "Centralized first-price clearing for budget-constrained two-sided multi-item markets," in *2026 39th Conference of Open Innovations Association (FRUCT)*. IEEE, 2026, pp. 214–221.

[12] T. Sandholm, "Algorithm for optimal winner determination in combinatorial auctions," *Artificial Intelligence*, vol. 135, no. 1–2, pp. 1–54, 2002.

[13] C. Boutilier and H. H. Hoos, "Bidding languages for combinatorial auctions," in *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 2*, ser. IJCAI'01. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2001, pp. 1211–1217.

[14] J. Cerquides, U. Endriss, A. Giovannucci, and J. A. Rodríguez-Aguilar, "Bidding languages and winner determination for mixed multi-unit combinatorial auctions," in *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, ser. IJCAI-07, 2007, pp. 1221–1226.

[15] M. H. Rothkopf, A. Pekeč, and R. M. Harstad, "Computationally manageable combinational auctions," *Management Science*, vol. 44, no. 8, pp. 1131–1147, 1998.

[16] T. Sandholm, S. Suri, A. Gilpin, and D. Levine, "Winner determination in combinatorial auction generalizations," in *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, ser. AAMAS '02. New York, NY, USA: Association for Computing Machinery, 2002, pp. 69–76.

[17] L. M. Ausubel and P. R. Milgrom, "Ascending auctions with package bidding," *The B.E. Journal of Theoretical Economics*, vol. 1, no. 1, 2002, art. no. 1.

[18] G. Goel, V. Mirrokni, and R. Paes Leme, "Polyhedral clinching auctions and the AdWords polytope," *J. ACM*, vol. 62, no. 3, pp. 18:1–18:27, Jun. 2015.

[19] H. Hirai and R. Sato, "Polyhedral clinching auctions for two-sided markets," *Mathematics of Operations Research*, vol. 47, no. 1, pp. 259–285, 2022.

[20] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. San Francisco, CA, USA: W. H. Freeman, 1979.

[21] H. Kellerer, U. Pferschy, and D. Pisinger, *Knapsack Problems*. Berlin, Heidelberg: Springer, 2004.