Centralized First-Price Clearing for Budget-Constrained Two-Sided Multi-Item Markets

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

John M. Ennis
NeoSwap Inc.
Midlothian, VA, USA
john.ennis@neoswap.ai

***Abstract*—Many multi-item platforms clear each item via an independent per-item auction. When buyers have hard budgets and per-buyer item caps, and sellers impose item-specific reserve prices, such parallel clearing ignores cross-item feasibility and can leave substantial gains from trade unrealized. This paper studies Smart Auction, a centralized clearing framework that retains first-price payment semantics while enforcing reserves, item caps, and ex post budget feasibility. We formulate allocation as a mixed-integer program that maximizes reported generalized surplus and implement a pay-as-bid mechanism (PAB) that charges winners their bids. To improve scalability, we propose a generalized surplus–first price mechanism (GSFP) that sorts buyer–item edges by surplus (bid minus reserve) and constructs a feasible allocation under caps and budgets in $O(E^+ \log E^+)$ time, where E^+ is the number of reserve-passing bid edges. In simulations over a grid of 4,320 markets per bidding regime (500–2,000 items, varying reserve and budget levels), GSFP improves welfare by about 35% and revenue by about 21% relative to a parallel first-price auction baseline, while closely matching PAB (less than 2.8% welfare and 1.7% revenue loss). GSFP runs roughly 70x faster than PAB for the tested scenarios, offering a practical welfare–revenue–runtime trade-off for constrained multi-item platforms.**

***Index Terms*—Multi-item auctions, budget constraints, reserve prices, first-price mechanisms, centralized clearing, greedy algorithms**

I. Introduction

Auctions provide a transparent way to allocate scarce resources and to discover prices, and they are used in many electronic markets today [1, 2]. In a large class of platforms, however, the traded objects are not a single homogeneous good but hundreds or thousands of heterogeneous, indivisible items offered simultaneously. It is common to clear such markets using parallel item-by-item auctions (e.g., position auctions in online advertising) [3, 4]. Parallel clearing is appealing for its simplicity and speed, but it ignores cross-item coupling created by practical feasibility constraints.

This paper focuses on two constraints that are central in multi-item platforms: (i) buyers often face *hard* spending limits (budgets) and *item caps* that restrict how many items they may win, and (ii) sellers (or the platform on their behalf) impose *item-specific reserve prices* that define minimum acceptable payments. These constraints interact across items. Allocating one item to a buyer reduces the set of other items that can be feasibly allocated to the same buyer, and conservative bid-submission heuristics can discard high-surplus trades ex ante. As a result, naive parallel clearing can leave substantial gains from trade unrealized when budgets and caps bind.

We study *Smart Auction*[1], a centralized clearing framework for two-sided multi-item markets with reserves, budgets, and item caps. Given sealed bids, Smart Auction computes a globally feasible allocation that maximizes reported generalized surplus (bid minus reserve) subject to single-assignment, cap, and budget constraints. This paper deliberately restricts attention to first-price payment semantics, because they allow ex post budget feasibility to be enforced directly via bid-based constraints.

Within Smart Auction, we implement two first-price mechanisms. *Pay-as-Bid* (PAB) solves a mixed-integer program (MIP) that maximizes reported generalized surplus and charges each winner her bid. *Generalized Surplus–First Price* (GSFP) replaces the MIP solve with a surplus-sorted greedy algorithm that runs in $O(E^+ \log E^+)$ time in the number of reserve-passing bid edges E^+, while maintaining the same feasibility checks. To quantify the value of centralized clearing, we compare both Smart Auction mechanisms against a *parallel-auction baseline* (Parallel-FP) that simulates independent per-item first-price auctions under a conservative, budget/cap-aware bid-submission rule.

Contributions:

(i) *A feasibility-constrained multi-item clearing model:* We formalize a two-sided, multi-item auction environment with item reserves, buyer budgets, and item caps, and express the generalized-surplus allocation as a compact MIP (Section III).

(ii) *Two first-price Smart Auction mechanisms:* We implement (i) PAB, an exact MIP-based allocation with first-price payments, and (ii) GSFP, a greedy surplus-first allocation rule that approximates the same objective with solver-free scalability (Section IV).

(iii) *Empirical comparison to parallel auctions:* Over 4,320 simulated markets per bidding regime (truthful bidding and fixed shading) with 500–2,000 items and varying reserve and budget levels, centralized Smart Auction

[1]The name *Smart Auction* refers to globally coordinated, constraint-aware clearing across items. It does not imply the use of machine learning.

clearing improves welfare and revenue substantially relative to Parallel-FP, while GSFP closely matches PAB at near-baseline runtime (Sections VI–VII).

II. Related Work

Our work connects centralized multi-item clearing, budget-feasible mechanism design, and practical auction formats used in large-scale platforms.

A. Multi-item auctions and welfare maximization

Classic results characterize truthful and efficient mechanisms in simpler settings, such as the second-price auction [5] and the Vickrey–Clarke–Groves (VCG) family for welfare maximization [6, 7]. In multi-item environments, full combinatorial auctions and exchanges generalize these ideas but can become computationally demanding as the number of items grows [8]. In contrast, many deployed systems rely on parallel per-item auctions due to scalability and implementation simplicity [3, 4].

B. Budgets, feasibility, and impossibility phenomena

Hard budget constraints fundamentally change incentives and feasibility. Even in bilateral trade, one cannot simultaneously satisfy incentive compatibility, individual rationality, budget balance, and efficiency [9] which is also the case in double auctions [10]. In richer multi-parameter environments, private or hard budgets lead to additional impossibility results and motivate approximate or restricted designs [11, 12]. Empirically and in platform practice, budgets are often treated as hard payment constraints, which shifts attention from fully truthful design to ex post feasible clearing rules.

C. Mechanisms under budgets and complex constraints

A substantial literature studies auctions with budget-constrained bidders, including models of financially constrained bidding [13] and algorithmic approaches that treat budgets as feasibility constraints [14]. Clinching-style mechanisms and their extensions handle ability-to-pay constraints in polyhedral feasibility environments [15, 16], and recent work extends clinching ideas to two-sided settings [17]. Our focus differs. We consider indivisible items with item-specific reserves and per-buyer item caps, and we study first-price instantiations of a centralized clearing rule that can be implemented either exactly (via MIP) or approximately (via a greedy rule).

D. Approximation, greedy allocation, and posted prices

When exact optimization is costly, greedy and posted-price mechanisms provide scalable alternatives, sometimes with formal approximation guarantees in restricted domains [18–21]. GSFP follows this approach. It uses surplus-sorted greedy selection to track the generalized-surplus objective while enforcing budgets and caps directly.

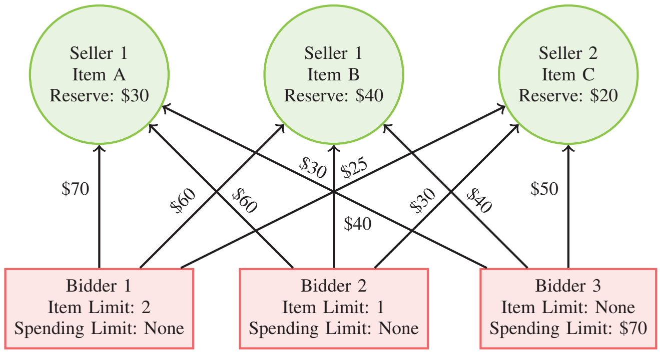

Fig. 1. An illustrative multi-item auction scenario with two sellers and three bidders. Arrow labels indicate bids.

E. Application context

Smart Auction is motivated by resale and exchange platforms in which participants face practical exposure constraints and sellers impose item-level minimum acceptable payments. It is also related to our earlier market models for circular-economy exchanges and resale platforms [22, 23], but the present work focuses specifically on centralized auction clearing and first-price payment semantics under hard feasibility constraints.

III. Model and Allocation Problem

We study a sealed-bid, two-sided multi-item market in which each item has a (seller-side) reserve price and each buyer faces hard feasibility constraints (a spending budget and an item cap). The key design question is how to clear the market when feasibility couples decisions across items.

A. Scenario description

To make the constraints concrete, consider the scenario in Fig. 1. There are three indivisible items to be auctioned. Seller 1 offers Items A and B with reservation prices $30 and $40, respectively, and Seller 2 offers Item C with a reservation price of $20.

There are three bidders. Bidder 1 submits bids of $70, $60, and $25 on Items A, B, and C, respectively, and may win up to two items (with no spending limit). Bidder 2 submits bids of $60, $40, and $30 on Items A, B, and C, respectively, and may win only one item. Bidder 3 submits bids of $30, $40, and $50 on Items A, B, and C, respectively, declares a spending limit of $70, and places no limit on the number of items won.

This scenario illustrates the cross-item coupling that motivates centralized clearing. For example, Bidder 3 can afford at most a subset of her bids, and Bidder 2 can win at most one item even if she is the top bidder on multiple items. A centralized rule can account for these constraints jointly when selecting the set of trades.

B. Items, sellers, buyers, bids, and constraints

Let $\mathcal{S}$ denote the set of sellers and $\mathcal{I}$ denote the set of items. Each item $i \in \mathcal{I}$ belongs to a seller $s(i) \in \mathcal{S}$ and is offered with a reserve price $r_i \geq 0$. Let $\mathcal{B}$ denote the set of buyers. Buyer $j \in \mathcal{B}$ has: (i) a spending limit (budget) $s_j \in \mathbb{R}_{\geq 0} \cup \{+\infty\}$ and (ii) an item cap $k_j \in \mathbb{Z}_{\geq 0} \cup \{+\infty\}$ limiting the number of items she may win. A missing limit ("None" in Fig. 1) is interpreted as $+\infty$. Equivalently, the corresponding constraint is dropped.

Each buyer may be interested in only a subset of items. We model bid eligibility via a set of buyer–item edges $\mathcal{E} \subseteq \mathcal{B} \times \mathcal{I}$. For each $(j,i) \in \mathcal{E}$, buyer j submits a sealed bid $b_{ji} \geq 0$ for item i. For each eligible pair (j,i) we also denote by v_{ji} buyer j's private value for item i.

An item i can be sold to at most one buyer, and each buyer is subject to both the cap and the budget constraints. Unsold items are explicitly allowed. If no reserve-passing feasible assignment is selected for item i, then i remains unallocated. Throughout this paper, we enforce budgets against the sum of winning bids. This convention is exact for first-price rules (winners pay their bid) and provides ex post budget feasibility by construction. If $s_j = +\infty$ (no spending limit), the budget constraint is redundant.

C. Reported generalized surplus and a compact MIP

Centralized Smart Auction clearing selects a feasible allocation that prioritizes trades with large reported surplus relative to reserves. For each eligible pair $(j,i) \in \mathcal{E}$, define the reported surplus as $\Delta_{ji} = b_{ji} - r_i$. Since trades below reserve are never desirable, we restrict attention to reserve-passing edges $\mathcal{E}^+ = \{(j,i) \in \mathcal{E} : b_{ji} \geq r_i\}$.

Let decision variables $x_{ji} \in \{0,1\}$ indicate whether buyer j is assigned item i. The generalized-surplus allocation problem can be written as the following MIP:

$$\max_x \quad \sum_{(j,i)\in\mathcal{E}^+} (b_{ji} - r_i)\, x_{ji} \tag{1}$$

$$\text{s.t.} \quad \sum_{j:(j,i)\in\mathcal{E}^+} x_{ji} \leq 1 \qquad \forall i \in \mathcal{I} \tag{2}$$

$$\sum_{i:(j,i)\in\mathcal{E}^+} x_{ji} \leq k_j \qquad \forall j \in \mathcal{B} \tag{3}$$

$$\sum_{i:(j,i)\in\mathcal{E}^+} b_{ji}\, x_{ji} \leq s_j \qquad \forall j \in \mathcal{B} \tag{4}$$

$$x_{ji} \in \{0,1\} \qquad \forall (j,i) \in \mathcal{E}^+. \tag{5}$$

Constraint (2) enforces single-assignment for each item. Constraint (3) enforces the item cap and (4) enforces budget feasibility on the sum of winning bids (consistent with first-price payments).

Problem (1)–(5) combines assignment constraints with per-buyer knapsack-like constraints induced by budgets, so it is NP-hard in general. Nevertheless, modern MIP solvers can solve moderate instances and provide a useful performance benchmark for scalable approximation rules such as GSFP.

D. Connections and computational hardness

The reported-surplus allocation problem in (1)–(5) can be viewed as a bipartite allocation problem with additional buyer-side knapsack-type constraints.

a) Polynomial special cases: If budgets are non-binding (or if we drop Constraint (4)) and each buyer has cap $k_j \in \{0,1\}$, the problem reduces to a maximum-weight bipartite matching / assignment problem, solvable in polynomial time (e.g., via the Hungarian algorithm) [24, 25]. More generally, without budgets the constraints define a b-matching polytope with integral extreme points, so linear relaxations can suffice in some regimes.

b) Hardness with budgets: Budget constraints fundamentally change the picture. Constraint (4) couples a buyer's accepted edges through the sum of winning bids. This introduces knapsack structure and makes the optimization NP-hard in general, even under simplified market structures [26, 27]. This computational barrier motivates scalable approximation rules such as GSFP, which trades optimality for speed while still enforcing the same ex post feasibility constraints.

IV. Smart Auction Mechanisms

Smart Auction consists of (i) a centralized allocation rule that enforces reserves, item caps, and budgets, and (ii) a payment rule. In this paper we use first-price payments throughout. Whenever buyer j wins item i, she pays $p_{ji} = b_{ji}$. Under first-price payments, enforcing budgets on the sum of winning bids (Constraint (4)) is sufficient for ex post budget feasibility.

The two mechanisms below differ only in how they compute a feasible allocation.

A. Pay-as-Bid (PAB)

PAB computes an allocation by solving the MIP in (1)–(5) to maximize reported generalized surplus subject to feasibility. The mechanism then applies first-price payments: if buyer j is allocated item i, she pays $p_{ji} = b_{ji}$. Because PAB enforces (4) on the sum of winning bids and charges winners their bids, payments are ex post budget feasible by construction. The primary drawback is computational. Solving a large MIP can dominate runtime as the number of items and bid edges grows (Fig. 3).

B. Generalized Surplus–First Price (GSFP)

GSFP is a scalable approximation to the same generalized-surplus objective. It replaces the MIP solve with a single greedy pass over reserve-passing edges, prioritizing trades with high surplus $b_{ji} - r_i$ while maintaining feasibility at each step.

Algorithm 1 matches the feasibility model in (2)–(4) during the greedy pass. An edge is accepted only if the item is unassigned, the buyer has remaining cap, and the buyer can afford the bid within the remaining budget. Because GSFP charges winners their bids, the budget check on S_j is exact for realized payments.

C. Runtime and complexity discussion

Let $E^+ = |\mathcal{E}^+|$ denote the number of reserve-passing bid edges. GSFP runs in $O(E^+ \log E^+)$ time:

(i) building $\mathcal{E}^+$ is linear in the input edge set;
(ii) sorting the reserve-passing edges dominates with $O(E^+ \log E^+)$, and

Algorithm 1 Generalized Surplus–First Price (GSFP)

```
Require: Reserves {r_i}_{i∈I}, bids {b_ji}_{(j,i)∈E}, item caps {k_j}, budgets {s_j}
Ensure: Allocation x_ji ∈ {0,1} and payments p_ji
Build E+ ← {(j,i) ∈ E : b_ji ≥ r_i} and surplus Δ_ji ← b_ji − r_i
Sort E+ by decreasing Δ_ji, and break ties by higher b_ji, then higher r_i, then a fixed deterministic index order (seller, item, buyer)
Initialize all items as unsold, and set remaining capacities K_j ← k_j and remaining budgets S_j ← s_j
for each (j,i) ∈ E+ in sorted order do
    if item i is unsold and K_j > 0 and S_j ≥ b_ji then
        Allocate i to j: set x_ji ← 1 and mark i as sold
        First-price payment: set p_ji ← b_ji
        Update capacities: K_j ← K_j − 1, S_j ← S_j − b_ji
    end if
end for
```

(iii) the subsequent feasibility scan processes each edge once in $O(1)$ time per edge.

In contrast, PAB requires solving an NP-hard MIP in the worst case. In our experiments, this computational gap is reflected in the runtime ratios in Table II and the scaling behavior in Fig. 3.

Note that PAB is exact with respect to the optimization problem in (1)–(5). When the MIP solver terminates at proven optimality (as in our experiments), the returned allocation maximizes reported generalized surplus over the stated feasible region. GSFP is not claimed to be globally optimal. It is a greedy heuristic for the same objective that preserves the same reserve, cap, and budget feasibility checks.

V. Economic and Algorithmic Properties

Because all mechanisms in this paper use first-price payments, they share several ex post feasibility properties. At the same time, incentive compatibility is inherently limited in hard-budget domains.

A. Ex post feasibility and balance

a) Reserve compatibility and seller individual rationality: All mechanisms trade only on reserve-passing edges $(j,i) \in \mathcal{E}^+$, so every sold item satisfies $b_{ji} \geq r_i$. Under first-price payments ($p_{ji} = b_{ji}$), each seller receives at least the reserve price for each sold item in gross terms, implying ex post individual rationality for sellers when reserves are interpreted as minimum acceptable gross payments. (If a platform commission is charged as a share of payments, seller IR can be enforced either by setting reserves net of commission or by interpreting r_i as a gross-floor. In our experiments the auctioneer share ratio is set to zero.)

b) Budget feasibility: PAB and GSFP enforce the budget constraint on the sum of winning bids directly (either via the MIP constraint (4) or via the greedy check $S_j \geq b_{ji}$ in Algorithm 1). Since realized payments equal winning bids, budgets are satisfied ex post by construction.

c) No positive transfers and budget balance: All payments are nonnegative (winners pay bids, losers pay zero), so the mechanisms satisfy no positive transfers. The mechanisms are also weakly budget balanced. Total buyer payments are redistributed to sellers (and, if applicable, to the platform as commission), so the mechanism never runs a deficit. When the auctioneer share ratio is zero, total buyer payments equal total seller revenues, giving strong budget balance in the standard sense.

B. Reported-surplus objective and individual rationality

a) Reported surplus and welfare gain: Define the *reported generalized surplus* of an allocation x as

$$\Phi(x; b, r) \;=\; \sum_{j\in\mathcal{B}} \sum_{i\in\mathcal{I}} (b_{ji} - r_i)\, x_{ji}.$$

The MIP in (1)–(5) maximizes $\Phi(x; b, r)$ over the feasible region. Under truthful bidding ($b_{ji} = v_{ji}$), $\Phi(x; v, r)$ coincides with our welfare-gain metric $W = \sum_{(j,i)\in\mathcal{T}} (v_{ji} - r_i)$, so the PAB allocation is welfare-gain optimal within the model's feasibility constraints. GSFP is a greedy approximation to maximizing $\Phi(x; b, r)$, while Parallel-FP is a heuristic baseline that does not directly optimize the global objective, since it decomposes the problem into per-item auctions after a conservative submission step.

b) Buyer individual rationality under no-overbidding: First-price payments make buyer utility sensitive to overbidding. If buyer j does not overbid, i.e., $b_{ji} \leq v_{ji}$ for all items, then any first-price outcome satisfies

$$U_j \;=\; \sum_{i:\, x_{ji}=1} (v_{ji} - b_{ji}) \;\geq\; 0,$$

so buyers are ex post individually rational relative to their true values. In particular, under truthful bidding and first-price payments, winners have zero utility (as in Table I), while bid shading can increase buyer utility by lowering payments.

c) Determinism and tie-breaking: To make outcomes reproducible and avoid spurious differences due to tie resolution, both centralized mechanisms use deterministic tie-breaking. GSFP sorts by (surplus, bid, reserve) and then by a fixed index order (seller, item, buyer), and PAB uses lexicographic objectives (maximize reported surplus, then total winning bid volume, then the number of trades).

C. Incentives

First-price mechanisms are not dominant-strategy incentive compatible (DSIC) in general. Bidders may strategically shade bids, and the outcome depends on beliefs and market thickness. More broadly, hard budgets create well-known tensions between incentive compatibility, individual rationality, and efficiency/balance goals. Even in simpler environments, it can be impossible to simultaneously achieve DSIC, individual rationality, Pareto efficiency, and budget balance [9, 10]. In richer multi-item domains with private budgets, deterministic DSIC mechanisms face additional impossibilities [11, 12].

Motivated by these limitations, this paper isolates the allocation and feasibility question: how much welfare and revenue can be recovered by moving from parallel item-by-item first-price clearing to centralized feasible clearing, and how close can a simple greedy allocation (GSFP) come to an exact surplus-maximizing benchmark (PAB) at practical runtimes.

VI. Experimental Setup

We evaluate Smart Auction's centralized first-price clearing using simulated markets generated by our instance generator. We compare two Smart Auction mechanisms, (PAB) and (GSFP), against a baseline that simulates parallel per-item first-price auctions under budgets and caps (Parallel-FP).

A. Baseline: Parallel Auctions (Parallel-FP Simulation)

Parallel-FP is a two-stage simulation intended to mimic item-by-item clearing when buyers face exposure constraints. It is used as the baseline for comparing outcomes of Smart Auction.

Stage 1 (bid submission under exposure constraints): Each buyer sorts her bids in decreasing order and iterates through this list, submitting a bid if doing so would not exceed: (i) the item cap k_j and (ii) the spending limit s_j when computed as the sum of submitted bid prices. Bids that would exceed the remaining budget are skipped, and the buyer continues with lower bids until reaching the item cap or exhausting her bid list. This submission rule is conservative. The sum of submitted bids upper-bounds any realized first-price payment, since only a subset of submitted bids can win.

Stage 2 (independent first-price clearing): For each item i, the auctioneer considers only submitted bids that meet the reserve ($b_{ji} \geq r_i$), selects the highest bid (with deterministic tie-breaking), and charges the winner her bid.

Pseudocode for parallel auction simulation is given in Algorithm 2.

Algorithm 2 Parallel First-Price baseline (Parallel-FP)

Require: Bids $\{b_{ji}\}$, reserves $\{r_i\}$, budgets $\{s_j\}$, caps $\{k_j\}$

```
Stage 1: cap/budget-aware bid submission
E_sub ← ∅
for each buyer j ∈ B do
    K ← k_j {remaining item cap}
    B ← s_j {remaining budget for submitted bids}
    for buyer j's bids (i, b_ji) in descending order do
        if K > 0 and B ≥ b_ji then
            E_sub ← E_sub ∪ {(j, i)}
            K ← K − 1;  B ← B − b_ji
        end if
    end for
end for
Stage 2: independent per-item first-price clearing
for each item i ∈ I do
    Choose j* ∈ arg max{b_ji : (j, i) ∈ E_sub, b_ji ≥ r_i}
    (ties broken deterministically)
    if such j* exists then
        Allocate i to j* and set payment p_{j*i} ← b_{j*i}
    end if
end for
```

B. Instance generator

Each instance contains $|\mathcal{I}|$ heterogeneous items with item-specific reserve prices. Items are grouped into sellers holding 1–5 items each. For each item i, an item quality score q_i is drawn i.i.d. from a lognormal distribution (log-mean 0, log-s.d. 0.5). The reserve price is then set as

$$r_i = \lfloor s_r \cdot (10\, q_i) \rfloor\,, \tag{6}$$

where $s_r \in \{0.5, 1.0, 2.0\}$ is the reserve scale parameter.

Buyers are heterogeneous. Each buyer j draws a type θ_j i.i.d. from the same lognormal distribution and an item cap k_j uniformly from $\{1, 2, 3, 4\}$. Buyer–item interest is sparse. For each buyer j, the number of items they are interested in is drawn as

$$L_j \sim \text{Poisson}(\lambda_B), \tag{7}$$

then truncated to satisfy $L_j \geq k_j$ and $L_j \leq |\mathcal{I}|$, and L_j distinct items are sampled uniformly without replacement. For each interest edge (j, i), we draw ε_{ji} lognormal and define the (true) private value

$$v_{ji} = \lfloor r_i + \lambda_v\, q_i\, \theta_j\, \varepsilon_{ji} \rfloor\,, \tag{8}$$

where $\lambda_v \in \{1.0, 2.0, 4.0\}$ is margin scale parameter. Values are clipped so that $v_{ji} \geq r_i$.

Budgets are exogenous and derived from values. For each buyer j, we consider only viable edges with strictly positive margin, i.e., items with $v_{ji} > r_i$. Let T_j denote the sum of the $\min\{k_j, |\mathcal{I}_j^+|\}$ largest values among these viable items, where $\mathcal{I}_j^+ = \{i : (j, i) \in \mathcal{E},\, v_{ji} > r_i\}$. Let $r_j^\star$ denote the reserve of j's favorite viable item (highest v_{ji} within $\mathcal{I}_j^+$). If $\mathcal{I}_j^+ = \varnothing$, we set $s_j = 0$. Otherwise, a buyer-specific budget ratio β_j is drawn from a Beta distribution with mean $\mu = \beta_{\text{budget}} \in \{0.1, 0.3, 0.7\}$ and fixed concentration $\kappa = 15$ (i.e., $\beta_j \sim \text{Beta}(\mu\kappa, (1-\mu)\kappa)$). The spending limit is then interpolated as

$$s_j = \text{round}\big((1 - \beta_j)\, r_j^\star + \beta_j\, T_j\big)\,. \tag{9}$$

In our experiments, the generator computes budgets from the unshaded values (i.e., with no shading for budget construction) and then the same budgets are reused across all bidding regimes and clearing rules.

C. Evaluation grid and bidding regimes

We evaluate a full factorial grid over

$$|\mathcal{I}| \in \{500, 1000, 1500, 2000\}, \quad \tau \in \{1, 3\}, \quad |\mathcal{B}| = \tau|\mathcal{I}|,$$

$$\beta_{\text{budget}} \in \{0.1, 0.3, 0.7\},\ s_r \in \{0.5, 1.0, 2.0\},$$

$$\lambda_v \in \{1.0, 2.0, 4.0\},\ \lambda_B \in \{5, 10\}.$$

For each parameter setting we generate 10 instances, yielding

$$N = 4 \cdot 2 \cdot 3 \cdot 3 \cdot 3 \cdot 2 \cdot 10 = 4320$$

markets per bidding regime.

We run two bidding regimes using a constant shading parameter α for all buyers and all clearing rules: (i) truthful bidding, $\alpha = 1.0$, and (ii) fixed shading, $\alpha = 0.85$. Given a true value v_{ji} and reserve r_i, the submitted bid is

$$b_{ji} = \lfloor \max\{r_i,\ r_i + \alpha\,(v_{ji} - r_i)\} \rfloor\,. \quad (10)$$

Budgets $\{s_j\}$ and caps $\{k_j\}$ are identical across clearing rules within an instance.

D. Metrics

Outcomes are computed using the true values $\{v_{ji}\}$ and reserves $\{r_i\}$ written by the generator. Let $\mathcal{T}$ denote the set of completed trades (allocated buyer–item pairs) and let p_{ji} denote the realized payment (first-price, so $p_{ji} = b_{ji}$ when trade occurs). We report:

$$W = \sum_{(j,i)\in\mathcal{T}} (v_{ji} - r_i) \quad \text{(true welfare / gains from trade),} \quad (11)$$

$$U = \sum_{(j,i)\in\mathcal{T}} (v_{ji} - p_{ji}) \quad \text{(buyers' true utility),} \quad (12)$$

$$S = \sum_{(j,i)\in\mathcal{T}} (p_{ji} - r_i) \quad \text{(gross sellers' surplus above reserves),} \quad (13)$$

$$R = \sum_{(j,i)\in\mathcal{T}} p_{ji} \quad \text{(revenue / total buyer payments).} \quad (14)$$

With first-price payments and zero auctioneer share (as in our experiments), welfare decomposes as $W = S + U$ at the instance level (Table I reports aggregated means, so W may differ slightly from $S + U$ after rounding). We also report the item trade rate as $|\mathcal{T}|/|\mathcal{I}|$ and the end-to-end runtime in seconds as recorded by the solver, summarized in Table II and visualized in Fig. 3.

VII. Experimental Results

We report results under truthful bidding ($\alpha = 1.0$) and fixed shading ($\alpha = 0.85$), using the same instance grid and seeds for all clearing rules. Our main comparisons focus on

(i) absolute performance aggregates across the grid (Table I),
(ii) paired improvements and losses computed instance-by-instance (Table II), and
(iii) (iii) diagnostic plots that explain where the gains originate (Figs. 2–7).

A. Aggregate outcomes across the grid

Table I reports mean and standard deviation across the full grid ($N = 4320$ instances per bidding regime). The following patterns are observed:

(i) *Centralized clearing increases welfare and trade:* Under truthful bidding, GSFP and PAB raise mean welfare gain W from 5843.9 (Parallel-FP) to 7617.3 and 7806.5, respectively. The higher welfare is accompanied by a higher item trade rate. The baseline sells about 82% of items on average, while GSFP and PAB sell about 98%–99%. With fixed shading, the same pattern holds. Both centralized mechanisms sell nearly all items and increase welfare relative to Parallel-FP.
(ii) *First-price shading primarily redistributes surplus:* Under truthful bidding, first-price payments imply that realized buyer utility is essentially zero on average ($U = 0$) and all welfare accrues as seller surplus ($S = W$). Under fixed shading ($\alpha = 0.85$), buyer payments fall below true values and buyer utility becomes positive. Table I shows this shift clearly. For each clearing rule, U rises under shading while S falls, with total welfare W changing more modestly. Fig. 2 provides a per-item visualization of this decomposition and the corresponding reduction in buyer payments per item.

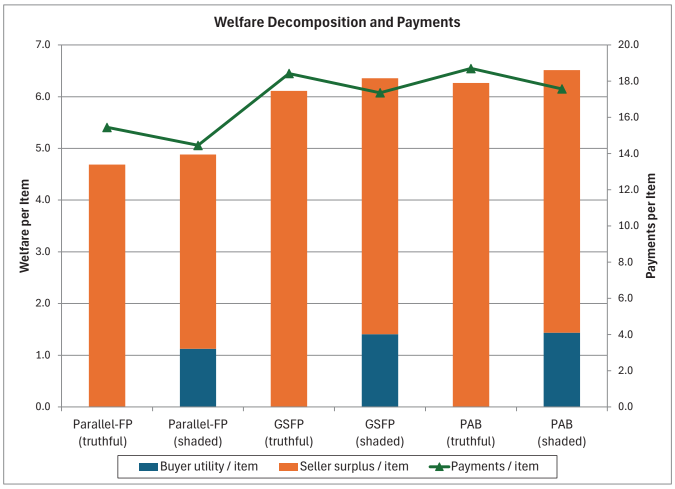

Fig. 2. Welfare decomposition and buyer payments for truthful bidding $\alpha = 1.0$ and fixed shading $\alpha = 0.85$. Stacked bars show per-item buyer utility and seller surplus. The overlaid line (secondary axis) shows per-item total buyer payments.

B. Paired comparisons and robustness

Aggregate means can mask heterogeneity, so Table II reports paired, instance-by-instance comparisons. The welfare and revenue improvements of GSFP and PAB over Parallel-FP are large and consistent. GSFP improves welfare by 35.9% (truthful) and 35.1% (shaded), and revenue by 20.4% (truthful) and 21.4% (shaded). Notably, GSFP beats Parallel-FP in 100% of instances for both welfare and revenue, indicating that the improvement is not driven by a small subset of markets.

The table also quantifies how close GSFP comes to the MIP benchmark. GSFP's mean welfare loss versus PAB is 2.8% in both bidding regimes, and its mean revenue loss is 1.7% (truthful) or 1.4% (shaded). Thus, across this grid, GSFP captures nearly all of the welfare and revenue gains delivered by MIP-based pay-as-bid clearing.

C. Runtime and scalability

The practical appeal of GSFP is that its gains do not come with a large computational penalty. Table II shows that the median runtime ratio GSFP/Parallel-FP is only $1.07\times$ (geometric mean approximately $1.10\times$), while PAB is roughly

TABLE I. AGGREGATE OUTCOMES ACROSS THE EVALUATION GRID (MEAN $\pm$ SD ACROSS THE $N = 4320$ INSTANCES IN EACH BIDDING REGIME). TRADE RATE IS THE FRACTION OF ITEMS SOLD. WELFARE DECOMPOSES AS $W = S + U$ (UP TO ROUNDING IN THE REPORTED MEANS).

Bidding	Clearing rule	Welfare Gain W	Sellers Surplus S	Buyers Utility U	Revenue	Item Trade Rate	Runtime (s)
Truthful bidding ($\alpha = 1.00$)	parallel-fp	5843.9 ± 5198.4	5843.9 ± 5198.4	0.0 ± 0.0	19277.8 ± 13212.5	$82\% \pm 14\%$	0.03 ± 0.03
	gsfp	7617.3 ± 6404.3	7617.3 ± 6404.3	0.0 ± 0.0	22988.1 ± 15466.3	$98\% \pm 5\%$	0.04 ± 0.03
	pab	7806.5 ± 6515.5	7806.5 ± 6515.5	0.0 ± 0.0	23333.5 ± 15595.8	$99\% \pm 3\%$	3.06 ± 2.57
Fixed shading ($\alpha = 0.85$)	parallel-fp	6085.6 ± 5435.1	4682.9 ± 4463.8	1402.7 ± 990.4	18037.1 ± 12505.0	$80\% \pm 14\%$	0.04 ± 0.03
	gsfp	7923.8 ± 6706.1	6165.4 ± 5551.9	1758.5 ± 1176.8	21666.2 ± 14709.3	$98\% \pm 3\%$	0.04 ± 0.03
	pab	8116.9 ± 6812.0	6324.8 ± 5648.7	1792.1 ± 1186.9	21928.5 ± 14819.2	$99\% \pm 2\%$	3.02 ± 2.75

TABLE II
PAIRED COMPARISON SUMMARY OVER $N = 4320$ INSTANCES. WELFARE AND REVENUE STATISTICS ARE IN %, RUNTIME STATISTICS ARE MULTIPLICATIVE FACTORS ($\times$). GSFP WELFARE/REVENUE LOSS VS PAB IS DEFINED AS $((W/R)_{\text{PAB}} - (W/R)_{\text{GSFP}})/(W/R)_{\text{PAB}}$.

Metric	Truthful bidding $\alpha = 1.0$	Fixed shading $\alpha = 0.85$
Welfare		
GSFP welfare gain vs Parallel-FP (mean, %)	35.9	35.1
PAB welfare gain vs Parallel-FP (mean, %)	40.0	39.2
GSFP welfare loss vs PAB (mean, %)	2.8	2.8
Share of instances where GSFP beats Parallel-FP (welfare, %)	100.0	100.0
Revenue		
GSFP revenue gain vs Parallel-FP (mean, %)	20.4	21.4
PAB revenue gain vs Parallel-FP (mean, %)	22.6	23.2
GSFP revenue loss vs PAB (mean, %)	1.7	1.4
Share of instances where GSFP beats Parallel-FP (revenue, %)	100.0	100.0
Runtime		
Runtime ratio PAB/GSFP (median, $\times$)	70.22	67.71
Runtime ratio PAB/GSFP (geometric mean, $\times$)	75.38	72.36
Runtime ratio GSFP/Parallel-FP (median, $\times$)	1.07	1.07
Runtime ratio GSFP/Parallel-FP (geometric mean, $\times$)	1.10	1.09

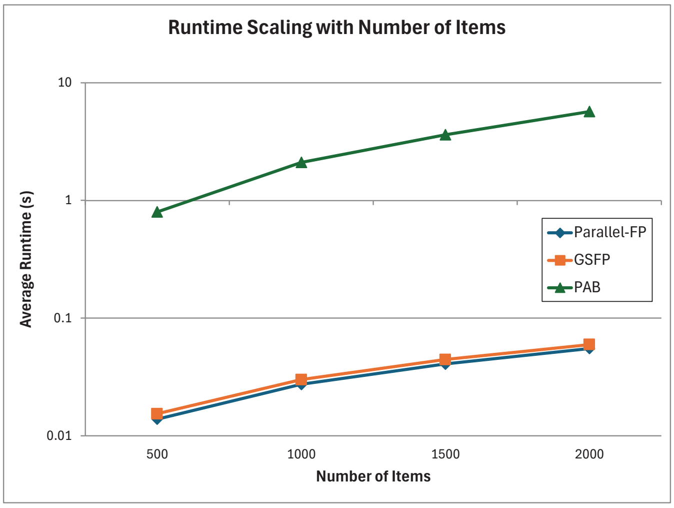

Fig. 3. Runtime scaling with the number of items for truthful bidding ($\alpha = 1.0$). The plot reports the average end-to-end runtime (seconds, logarithmic y-axis) aggregated over all other parameter settings and evaluation seeds. GSFP closely tracks the Parallel-FP baseline while PAB is substantially slower.

70$\times$ slower than GSFP on median. Fig. 3 further illustrates how runtime scales with the number of items. GSFP closely tracks Parallel-FP as the market grows from 500 to 2,000 items, whereas PAB becomes substantially slower. While we do not run a separate scalability experiment, the evaluation grid already spans up to 2,000 items, and the scaling plot aggregates over all other parameters and seeds.

D. *Where do the welfare and revenue gains come from?*

	Budget Ratio			
Reserve Scale	Low (0.1)	Medium (0.3)	High (0.7)	Overall
Low (0.5)	20.0%	25.0%	19.4%	**21.3%**
Medium (1)	31.4%	35.3%	24.9%	**29.9%**
High (2)	45.0%	45.6%	31.0%	**39.5%**
Overall	**32.9%**	**35.4%**	**25.0%**	

Min	Median	Max
19.4%	30.7%	45.6%

Fig. 4. Heatmap (truthful bidding, $\alpha = 1.0$) shows the percentage improvement in mean welfare gain of GSFP over the Parallel-FP baseline, aggregated by reserve scale and budget ratio. Each cell reports $(\overline{W}_{\text{GSFP}}/\overline{W}_{\text{Parallel-FP}}) - 1$, where each $\overline{W}$ is the mean welfare gain across instances in that bin. The Overall row/column aggregate across the corresponding dimension.

	Budget Ratio			
Reserve Scale	Low (0.1)	Medium (0.3)	High (0.7)	Overall
Low (0.5)	14.2%	18.7%	15.1%	**16.1%**
Medium (1)	19.0%	21.9%	16.2%	**18.9%**
High (2)	23.6%	23.0%	16.1%	**20.7%**
Overall	**20.5%**	**21.8%**	**15.9%**	

Min	Median	Max
14.2%	19.0%	23.6%

Fig. 5. Heatmap (truthful bidding, $\alpha = 1.0$) shows the percentage improvement in mean revenue (total buyer payments) of GSFP over the Parallel-FP baseline, aggregated by reserve scale and budget ratio. Each cell reports $(\overline{R}_{\text{GSFP}}/\overline{R}_{\text{Parallel-FP}}) - 1$, where each $\overline{R}$ is the mean total buyer payments across instances in that bin. The Overall row/column aggregate across the corresponding dimension.

The heatmaps in Figs. 4–7 summarize welfare and revenue improvements by reserve scale and budget ratio (using within-bin averages), complementing the instance-by-instance paired summaries in Table II. Improvements are generally largest when (i) reserve prices are high (making poor coordination more costly) and (ii) budgets are tight-to-moderate (increasing cross-item coupling), with the smallest gains when budgets are very loose. In these regimes, Parallel-FP's item-by-item clearing and submission heuristics more frequently block feasible high-surplus bundles. Centralized clearing remedies this by selecting a globally feasible set of trades that prioritizes high surplus edges, either exactly via the MIP (PAB) or approximately via the surplus-sorted greedy rule (GSFP). Across bins, the structure of gains is similar for GSFP and PAB, consistent with the small average performance gap between them in Table II.

	Budget Ratio			
Reserve Scale	Low (0.1)	Medium (0.3)	High (0.7)	Overall
Low (0.5)	23.7%	29.0%	21.8%	**24.5%**
Medium (1)	35.3%	39.3%	27.2%	**33.1%**
High (2)	49.1%	49.5%	33.1%	**42.7%**
Overall	**36.8%**	**39.4%**	**27.3%**	**33.6%**

Min	Median	Max
21.8%	33.4%	49.5%

Fig. 6. Heatmap (truthful bidding, $\alpha = 1.0$) shows the percentage improvement in mean welfare gain of PAB over the Parallel-FP baseline, aggregated by reserve scale and budget ratio. Each cell reports $(\overline{W}_{\text{PAB}}/\overline{W}_{\text{Parallel-FP}})-1$ using per-bin means. The Overall row/column aggregate across the corresponding dimension.

	Budget Ratio			
Reserve Scale	Low (0.1)	Medium (0.3)	High (0.7)	Overall
Low (0.5)	17.7%	21.4%	16.7%	**18.5%**
Medium (1)	21.7%	24.0%	17.3%	**20.8%**
High (2)	25.8%	24.6%	16.8%	**22.2%**
Overall	**23.1%**	**23.8%**	**16.9%**	

Min	Median	Max
16.7%	21.2%	25.8%

Fig. 7. Heatmap (truthful bidding, $\alpha = 1.0$) shows the percentage improvement in mean revenue (total buyer payments) of PAB over the Parallel-FP baseline, aggregated by reserve scale and budget ratio. Each cell reports $(\overline{R}_{\text{PAB}}/\overline{R}_{\text{Parallel-FP}})-1$ using per-bin means. The Overall row/column aggregate across the corresponding dimension.

VIII. Conclusion

We presented *Smart Auction*, a centralized clearing framework for two-sided multi-item markets with item-specific reserves, buyer budgets, and item caps, and we studied two practical instantiations that retain first-price payment semantics. Pay-as-Bid (PAB) computes the generalized-surplus-maximizing allocation via a mixed-integer program (MIP) and charges winners their bids, while Generalized Surplus–First Price (GSFP) uses a surplus-sorted greedy rule with the same feasibility checks.

Across a grid of 4,320 simulated markets with up to 2,000 items, both Smart Auction mechanisms substantially increase welfare, revenue, and trade rate relative to a parallel per-item first-price baseline (Parallel-FP). GSFP captures nearly all of the welfare and revenue gains of MIP-based PAB while running near baseline speed. In our experiments, PAB is typically on the order of $70\times$ slower than GSFP, whereas GSFP is only about $1.07\times$ slower than Parallel-FP.

A natural direction for future work is to study alternative pricing rules and strategic bidding models on top of the same feasibility-constrained clearing problem, including second-price or VCG-inspired payments, posted-price variants, and richer bidder behavior under exposure constraints.

References

[1] R. P. McAfee and J. McMillan, "Auctions and bidding," *Journal of Economic Literature*, vol. 25, no. 2, pp. 699–738, 1987.
[2] V. Krishna, *Auction Theory*. Academic Press, 2002.
[3] B. Edelman, M. Ostrovsky, and M. Schwarz, "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords," *American Economic Review*, vol. 97, no. 1, pp. 242–259, 2007.
[4] H. R. Varian, "Position auctions," *International Journal of Industrial Organization*, vol. 25, no. 6, pp. 1163–1178, 2007.
[5] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *Journal of Finance*, vol. 16, no. 1, pp. 8–37, 1961.
[6] E. H. Clarke, "Multipart pricing of public goods," *Public Choice*, vol. 11, pp. 17–33, 1971.
[7] T. Groves, "Incentives in teams," *Econometrica*, vol. 41, no. 4, pp. 617–631, 1973.
[8] D. C. Parkes, J. R. Kalagnanam, and M. Eso, "Achieving budget-balance with vickrey-based payment schemes in combinatorial exchanges," in *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, 2001, pp. 1161–1168.
[9] R. B. Myerson and M. A. Satterthwaite, "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, vol. 29, no. 2, pp. 265–281, 1983.
[10] R. P. McAfee, "A dominant strategy double auction," *Journal of Economic Theory*, vol. 56, no. 2, pp. 434–450, 1992.
[11] S. Dobzinski, R. Lavi, and N. Nisan, "Multi-unit auctions with budget limits," *Games and Economic Behavior*, vol. 74, no. 2, pp. 486–503, 2012.
[12] P. Le, "Pareto optimal budgeted combinatorial auctions," *Theoretical Economics*, vol. 13, no. 2, pp. 831–868, 2018.
[13] Y.-K. Che and I. Gale, "Standard auctions with financially constrained bidders," *Review of Economic Studies*, vol. 65, no. 1, pp. 1–21, 1998.
[14] A. Fiat, S. Leonardi, J. Saia, and P. Sankowski, "Single valued combinatorial auctions with budgets," in *Proceedings of the 12th ACM Conference on Electronic Commerce*, ser. EC '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 223–232.
[15] G. Goel, V. Mirrokni, and R. Paes Leme, "Clinching auctions beyond hard budget constraints," in *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, ser. EC '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 167–184.
[16] G. Goel, V. Mirrokni, and R. P. Leme, "Polyhedral clinching auctions and the adwords polytope," *J. ACM*, vol. 62, no. 3, Jun. 2015.
[17] H. Hirai and R. Sato, "Polyhedral clinching auctions for two-sided markets," *Math. Oper. Res.*, vol. 47, no. 1, p. 259–285, 2022.
[18] D. Lehmann, L. O'Callaghan, and Y. Shoham, "Truth revelation in approximately efficient combinatorial auctions," *Journal of the ACM*, vol. 49, no. 5, pp. 577–602, 2002.
[19] A. V. Goldberg, J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright, "Competitive auctions," *Games and Economic Behavior*, vol. 55, no. 2, pp. 242–269, 2006.
[20] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan, "Multi-parameter mechanism design and sequential posted pricing," in *Proceedings of the Forty-Second ACM Symposium on Theory of Computing (STOC '10)*. Association for Computing Machinery, 2010, pp. 311–320.
[21] M. Feldman, N. Gravin, and B. Lucier, "Combinatorial auctions via posted prices," in *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2015)*. Society for Industrial and Applied Mathematics, 2015, pp. 123–135.
[22] A. H. Özer, "A double auction based mathematical market model and heuristics for internet-based secondhand durable good markets," *Computers & Operations Research*, vol. 111, pp. 116–129, 2019.
[23] ——, "A fair, preference-based posted price resale e-market model and clearing heuristics for circular economy," *Applied Soft Computing*, vol. 106, p. 107308, 2021.
[24] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, pp. 83–97, 1955.
[25] J. Munkres, "Algorithms for the assignment and transportation problems," *Journal of the Society for Industrial and Applied Mathematics*, vol. 5, no. 1, pp. 32–38, 1957.
[26] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. San Francisco, CA: Freeman, 1979.
[27] H. Kellerer, U. Pferschy, and D. Pisinger, *Knapsack Problems*. Berlin, Germany: Springer, 2004.