Robust Double Auctions for Resource Allocation

Arthur Lazzaretti[1⋆], Charalampos Papamanthou[1,2], and Ismael Hishon-Rezaizadeh[2]

Yale University[1]
Lagrange Labs[2]

Abstract. In a zero-knowledge proof market, we have two sides. On one side, bidders with proofs of different sizes and some private value to have this proof computed. On the other side, we have distributors (also called sellers) which have computational power available to process the proofs by the bidders, and these distributors have a certain private cost to process these proofs (dependent on the size). More broadly, this setting applies to any online resource allocation where we have bidders who desire a certain amount of a resource and distributors that can provide this resource. In this work, we study how to devise double auctions for this setting which are truthful for users, weak group strategy proof, weak budget balanced, computationally efficient, and achieve a good approximation of the maximum welfare possible by the set of bids. We denote such auctions as *robust*.

1 Introduction

Consider the setting of a zero-knowledge proof market, where bidders are interested in having a zero-knowledge proof computed for a circuit they have (for a certain price) and distributors (also referred to as sellers) are willing to perform the proof computation for these bidders (for a certain price). In the middle, we have an auctioneer who is in charge of matching bidders and distributors as well as deciding the final prices—how much each matching bidder has to pay and how much each matching distributor is receiving. Clearly, there are incentives in this setting: the bidders want to have their proof computed for the least amount of money as possible, and the distributors want to maximize their revenue given their cycles of compute. How should the auctioneer decide which bidders to match with which distributors and at what prices?

In particular, our auction setup is the following: The market has n bidders interested in circuits of different complexities and sizes. Each bidder i has circuit of size k_i (measured in required cycles computation) and private value v_i for having the zero-knowledge proof for this circuit computed. The market also has t distributors, each distributor j with a different computation capacity (units of computation) u_j they can provide per round, and a cost per compute cycle c_j.

Our auction setup has a series of caveats. The bidders are only interested in a certain number of cycles of compute, corresponding to the number of cycles

⋆ Work done during author's time at Lagrange Labs.

needed to prove their circuit, and no less. In the literature, this has been coined as a "single-minded bidder" [6] (we will also call them knapsack bidders). Distributors do not have this same "all-or-nothing" constraint. They are happy to sell however many cycles of compute they have available, for the same price per cycle. Furthermore, since proofs can be (and most times need to be) distributed we do not impose any constraints on a bidder being assigned to one or multiple distributors. Due to the above, and in order to simplify the model, we in many places consider a distributor with u_j units of capacity to instead be u_j different distributors selling one cycle of compute at cost c_j. More generally, as modeled, this setting encompasses a double auction of any resource where the bidder is interested in only a certain quantity of this resource and no less. Other examples can include storage, general computation, ML training, ad space or time and others. (For a physical example, consider a chair auction, where bidders have different sized tables and either want chairs for their whole table, or no chairs at all.)

Related works on double auctions. Two-sided markets are used throughout the web for a myriad of different settings and goods, stocks and futures, currencies, ad sales [2], among others [15,24,9] and therefore double auctions have been very well studied. In particular, knapsack auctions have been studied in the single-sided auction setting [2,6] (multiple bidders, single distributor) but not in the double-auction setting. On the other hand, double auctions have been studied in the case where bidders are interested in a single item, or are interested in many items, but have a *diminishing* value per additional item of interest—as opposed to a value dictated by a knapsack constraint. Another line of work looks at bundles, when sellers are restricted to selling one item per commodity. (We give a more detailed overview of related work in Appendix A.)[1]

The problem we pose is to define a knapsack double auction mechanism which has the following properties:

- Truthfulness: The best strategy for auction participants is to report their private values to the auctioneer.
- (Weak) Group-Strategy Proofness: For any group of colluding participants, for at least one of the participants in the group, the best strategy is still to bid truthfully over participating in the collusion.
- Welfare-Maximization: Across all possible assignments that satisfy the knapsack constraint, you pick the ones with the largest sum of private values.
- Weak Budget-Balanced: The auctioneer does not lose money from facilitating the auction.
- Computational Efficiency: The mechanism runs in polynomial time in the size of the inputs.

These properties are standard notions considered for auction mechanisms. We define a double auction to be robust if it satisfies all these properties, with the caveat of the welfare-maximization being an approximate notion (we leave

[1] Recent work by Wang et al. [27] studies a similar setting but with important differences. We discuss this in a dedicated subsection.

the exact approximation to be defined in Section 5). This is because by the Myerson-Satterthwaite theorem [23], no auction mechanism can satisfy all of these properties at once (even without the weak group-strategy proofness requirement). Next, we consider existing mechanisms to check whether they may satisfy our needs.

Vickrey-Clarke-Groves (VCG) Knapsack Double Auction [26,11,13,7]. Applying the double-auction version of the classical Vickrey-Clarke-Groves (VCG) auction naively fails: It is not computationally efficient, not budget balanced [18] (although it is collusion-proof [8]). There has been some VCG-style approaches to this idea which either do not achieve budget balance or do not give a welfare analysis in the general setting. We expand on this in related work.

Pay-as-bid Double Auction. Another option would be an orderbook-style approach, where bidders and distributors report their cost and transact whenever there is a match in ranges. In this case, notice that the one to bid first can have an incentive to strategize based on expected demand and other bids it sees within the orderbook. In addition, the auction has discriminatory pricing (the same item can be sold for different prices). This makes the auction not truthful (this is because for certain bidder and distributor distributions, an individual bidder could increase their utility by bidding something other than their true private value).

1.1 Our contribution

This leads us to question whether we can devise a framework for knapsack double auction mechanisms that are *robust* as defined above. Unfortunately, we show that achieving a perfectly welfare-maximizing solution and a computationally efficient solution is impossible, since it is equivalent to the knapsack problem which is known to be NP. Therefore, we tweak our definition of robustness to allow for an approximate welfare-maximization.

Our contribution in this work is to show that there is a class of knapsack double auctions that satisfy the properties above (with the approximate welfare-maximization). We achieve this through a modular approach, inspired by the composition framework of Dütting et al. [12].

Dütting et al. [12] established a framework where the input is two ranking algorithms, for user bids and distributors bids respectively, a composition rule and a payment rule, and the output is a double auction mechanism. They show that for the scenario where multiple bidders and distributors interested in buying/selling a single item, if the ranking algorithms, composition rule and payment rule exhibited certain properties, then we could conclude that the resulting double auction mechanism was truthful, weak GSP, welfare-maximizing (to a certain extent) and WBB.

In our work, we follow the same approach, extending their framework to the knapsack setting, where bidders are interested in only a certain number of items (or no items) and distributors don't have this constraint. Extending the framework requires entirely new definitions, proofs and techniques. This general framework allows for more flexibility in devising what we call knapsack double auctions by encompassing most current knapsack solutions. Specifically, we don't show that a single specific mechanism satisfies the properties above, but instead show that the properties hold for a whole class of functions that satisfy our constraints. (The work of extending the composition framework to some cross-auction constraints was posed as an interesting open question by the work in [12]; this work answers that open question in establishing the feasibility of establishing a composition framework which can support cross-auction constraints.)

2 Preliminaries

In this section we work through definitions and some background that will be necessary when approaching our problem.

2.1 Double Auctions

We will study an asymmetric double auction setting. It is a two sided market with n bidders and t distributors . There is one item type for sale. Each distributor has u units to sell at a price c each[2], and each bidder is interested in either getting k units, or no units. We call this bidder a 'single minded bidder' or a 'knapsack bidder'. For the special case where $k = 1$ for every bidder, this degrades to the multi-bidder multi-distributor single item auction, which is covered [12]. In this setting (where $k = 1$ for every bidder) there have been other works that are robust [21].

A set of bidders and distributors is *feasible* if there are at least as many distributors as the the sum of the items desired by the set of bidders, $\sum_{i\in[N]} k_i$.

Each bidder i has a value v_i and size k_i, and each distributor j has a unit capacity u_j and a *cost per unit* c_j, all of these are bounded from above and below by some maximum and minimum value. We will denote by $\boldsymbol{b} = (\boldsymbol{v}, \boldsymbol{k})$ the value and size profile of each bidder. We will denote by $b_i = (v_i, k_i)$ the bid representing the value and size of the i-th bidder. We denote by $\boldsymbol{d}$ the cost profile of all distributors. We will denote by $d_j = (u_j, c_j)$ the unit capacity and cost per unit of distributor j.

We define a double auction mechanism as a tuple of two deterministic functions:

- The allocation rule $x(\cdot,\cdot)$: It takes in the inputs of bidders and distributors and outputs the set of *winning*[3] bidders and distributors. For every $i \in [n]$, we let $x_i(\boldsymbol{b}, \boldsymbol{d}) \in \{0, 1\}$ to denote whether bidder i was a winning bidder. Symmetrically, we define $x^j(\boldsymbol{b}, \boldsymbol{d})$ to denote a vector of size $m = \sum_{\ell\in[t]} u_\ell$ where each index denotes whether each unit of a certain distributor was allocated, in order.

[2] To simplify the exposition, where appropriate we will split a distributor into u single-unit distributors.

[3] Where we define winning bidders and distributors to be those who actually transact as a result of the auction.

– The payment rule $p(\cdot,\cdot)$: The payment rule takes in the same inputs as $x(\cdot,\cdot)$ and outputs the price each winning bidder is required to pay, and payment each distributor is entitled to receive. We will denote $p_i(\boldsymbol{b},\boldsymbol{d})$ to be the price to be paid by bidder i and $p^j(\boldsymbol{b},\boldsymbol{d})$ the payment to be received by distributor j for each unit (again p^j is of size m, we break each distributor up for the assignment). For any non-winning bidder/distributor units, this value is 0.

We can then define the welfare of bidder i from acquiring k_i units to be $x_i(\boldsymbol{b},\boldsymbol{d})\cdot v_i - p_i(\boldsymbol{b},\boldsymbol{d})$ (the price paid when a user is not selected is 0). Analogously we define the welfare of distributor j for selling a unit to be $p^j(\boldsymbol{b},\boldsymbol{d}) - x^j(\boldsymbol{b},\boldsymbol{d})\cdot c_j$. Finally, we can define the welfare of the auctioneer holding the auction. This is just the difference between the prices paid by the bidders and the payments made to the distributors, $\sum_{i\in[n]} p_i(\boldsymbol{b},\boldsymbol{d}) - \sum_{j\in[m]} p^j(\boldsymbol{b},\boldsymbol{d})$.

Now, we can define the *welfare of the auction mechanism* as the sum of the welfare of all participating parties. After summing and cancelling some terms, we get that:

$$\mathsf{W}(\boldsymbol{b},\boldsymbol{d}) = \sum_{i\in[n]} x_i(\boldsymbol{b},\boldsymbol{d})\cdot v_i - \sum_{j\in[m]} x^j(\boldsymbol{b},\boldsymbol{d})\cdot c_j\,.$$

Because we also account for the auctioneer's welfare, the welfare is actually independent of the payment rule. Intuitively this means that our welfare definition does not necessarily capture payment 'fairness' as in who should the excess welfare belong to. However, this will be (partially) covered by our truthfulness which holds for both sides.

2.2 Properties for Double Auction Mechanisms

We now define the properties we want double auction mechanisms to satisfy more formally.

Truthfulness: A double auction mechanism is *truthful* if for every bidder or distributor, reporting the private value (resp. the true cost) for the good(s) is a best strategy. Furthermore, any participating bidder or distributor cannot get negative utility.

Formally, this means that, for any bidder $i \in [n]$ with truthful input $b_i = (v_i, k_i)$, and for any b'_i:

$$x_i(\boldsymbol{b},\boldsymbol{d})\cdot v_i - p_i(\boldsymbol{b},\boldsymbol{d}) \geq x_i((b'_i,\boldsymbol{b}_{-i}),\boldsymbol{d})\cdot v_i - p_i((b'_i,\boldsymbol{b}_{-i}),\boldsymbol{d})\,.$$

Furthermore, for any bidder $i \in [n]$ that bids truthfully, $x_i(\boldsymbol{b},\boldsymbol{d})\cdot v_i - p_i(\boldsymbol{b},\boldsymbol{d}) \geq 0$. It is defined symmetrically for distributors.

Weak Group Strategy Proofness (WGSP): A double auction mechanism is WGSP if for any input bid and cost vectors $\boldsymbol{b},\boldsymbol{v}$ to the mechanism, for every set $A \subseteq (B \bigcup S)$, subset of the participating bidders and distributors, and every alternative reporting bid and/or cost inputs the participants in A can pick together, there is always at least one party in A who is no better off by participating in A than it is reporting truthfully. This means that a colluding set of parties cannot skew the auction so that all colluding parties get better utility from the auction.[4]

(Weak) Budget Balancing ((W)BB): A double auction mechanism is budget balanced if the auctioneer neither gains nor loses utility from conducting the auction. It is *weak* budget balanced if the auctioneer does not lose utility (in this case the auctioneer can gain utility).

Welfare Maximizing: A double auction mechanism is welfare maximizing if the welfare $W(\cdot,\cdot)$ output by the mechanism for any input sets $\boldsymbol{b}, \boldsymbol{d}$ is exactly $\mathsf{OPT}(\boldsymbol{b}, \boldsymbol{d})$, where we define OPT as follows:

$$\mathsf{OPT}\left(\boldsymbol{b} = (\boldsymbol{v}, \boldsymbol{k}), \boldsymbol{d} = (\boldsymbol{u}, \boldsymbol{c})\right) = \max_{B,S: B\subseteq[n], S\subseteq[m] \text{ s.t. } \sum_{i\in B} k_i \leq |S|} \left\{ \sum_{i\in B} v_i - \sum_{j\in S} c_j \right\}.$$

We will use this function $\mathsf{OPT}(\boldsymbol{b}, \boldsymbol{d})$ throughout to denote the maximum possible welfare achievable by a set of bidders and distributors. We say that a mechanism is $\epsilon-$welfare maximizing if for any $\boldsymbol{b}, \boldsymbol{d}$, $W(\boldsymbol{b}, \boldsymbol{d}) \geq \epsilon \cdot \mathsf{OPT}(\boldsymbol{b}, \boldsymbol{d})$.

Notice that given a blackbox algorithm to solve $\mathsf{OPT}(\boldsymbol{b}, \boldsymbol{d})$ for any $\boldsymbol{b}, \boldsymbol{d}$, we can construct a solver to the knapsack problem by assigning all costs equals 0 and picking the number of distributors equal the weight of the knapsack and assigning bidders with size equal value equal weight. This implies that computing $\mathsf{OPT}(\boldsymbol{b}, \boldsymbol{d})$ is at least as hard as knapsack and therefore not solvable in polynomial time.

Computationally Efficient: A double auction mechanism is computationally efficient if for any input $\boldsymbol{b}, \boldsymbol{d}$, both of its functions $x(\cdot,\cdot)$ and $p(\cdot,\cdot)$ can be computed in time polynomial in the size of the input.

2.3 Relevant Prior Work on Double Auctions and Compositions

We first outline a known characterization of double auctions with respect to truthfulness. For that, we will need two definitions.

Definition 21 (Monotone Allocation Rule [12]) *An allocation rule is $x(\cdot,\cdot)$ is monotone if for every input set $\boldsymbol{b}, \boldsymbol{d}$, every winning bidder (resp. distributor) who raises his value (resp. lowers his cost) while other bids and costs remain static is still a winner.*

This monotone allocation rule just states that bidders should never be removed from the winning set for bidding more (resp. distributors should not be removed from the winning set for bidding a lower cost.) Next we define a threshold payment rule.

[4] Group Strategy Proofness (without the weak) says no group can collude to make some party better off while the others don't lose anything. This cannot be satisfied by any double auction mechanism [12] so we don't consider it.

Definition 22 (Threshold Payments [12]) *The threshold payment for bidder i (resp. for distributor j), given inputs* $\boldsymbol{b}_{-i}, \boldsymbol{d}$ *(resp.* $\boldsymbol{b}, \boldsymbol{d}_{-j}$*) and a monotone allocation rule* $x(\cdot,\cdot)$*, are respectively:*

$$\inf_{v_i : x_i(\boldsymbol{b},\boldsymbol{d})=1} v_i \, , \sup_{c_j : x^j(\boldsymbol{b},\boldsymbol{d})=1} c_j \, .$$

What this is saying is that the threshold payment is the minimum value this winning bidder would need to bid to still be accepted. The intuition is symmetrical for distributors.

Theorem 21 (Truthful Double Auction Mechanisms [12]). *A double auction mechanism is truthful if and only if the allocation rule is monotone and the payment rule applies threshold payments.*

The theorem above assumes that participants with a value of 0 always receive or pay 0.

Ranking Algorithms for Compositions [12] The only definition we use (almost) as-is from [12] is the definition for ranking algorithms for single unit distributors. In their work, they use this definition to compose single unit distributors and single unit bidders. Although we use the same general framework, our definitions need changes in order to suit the new multi-unit singleminded bidders and multi-unit sellers, so we redefine them in the next section, along with new definitions of monotonicity and consistency, which are also properties defined previously but which require changes to make sense in the multi-unit setting.

3 A New Composition Framework

In this section, we establish our new composition framework which encompasses single minded bidders interested in multiple items. We will call these *asymmetric* compositions, which compose a bidder with a knapsack constraint on its value, with distributors that are interested in selling a single item.

3.1 Ranking Algorithms

We will use one-sided *ranking algorithms* for user and operator bids to be used for our composition.

Definition 31 (Knapsack Ranking) *A knapsack ranking algorithm* $r(\boldsymbol{b})$ *is a deterministic algorithm that takes in a vector of* n *single minded bids of the form* $b_i = (v_i, k_i)$*, where* v_i *is the value this bidder assigns to getting* k_i *items, and outputs an ordered list of tuples. We will denote* $r_i(\boldsymbol{b})$ *to be the* $i-th$ *element of the ordered output of* r*.*

In this definition, the ordering itself could be dependent of v_i, k_i, or both. We will see in the next sections the properties we will need from this ordering algorithm in order to achieve our desired properties.

Definition 32 (Cost Ranking) *A cost ranking algorithm $\hat{r}$ takes in a vector of distributor bids $\boldsymbol{d}$ with each bid j of the form $\boldsymbol{d_j} = (u_j, c_j)$, where u_j is the unit capacity and c_j is the cost per unit of distributor j. The output $\hat{r}(\boldsymbol{d})$ is a ordered list of tuples. We will denote $\hat{r}_j(\boldsymbol{d})$ to be the j-th available unit in $\hat{r}(\boldsymbol{d})$.*

Similar to above, we leave how exactly the order is picked open, however, specifically for the case of distributors, in our setting the only ordering that makes sense is to sort the distributors in increasing order by cost.

Notation. Notice that we have used a slightly different notation for indexing each ranking. Whereas we index the knapsack ranking by tuple, we index the cost ranking per unit. For example, if $\hat{r}(\boldsymbol{d}) = (5, 3), (2, 7)$ then $\hat{r}_4(\boldsymbol{d}) = 3$ (there are 5 units of cost 3 and 2 units of cost 7).

Definition 33 (Asymmetric Composition Rule) *An asymmetric composition rule for ranking algorithms for multi-unit single-minded bidders (resp. distributors) and simple distributors (resp. bidders) receives as input a knapsack user bid $b = (v, k)$ and a cost per item c and outputs either* 1 *(accept) or* 0 *(reject).*

We now define the composition rule which will be the one we use the most throughout the paper.

Definition 34 (t-Threshold Asymmetric Composition Rule) *The $t-$threshold composition rule takes in a bid $b = (v, k)$, the value attributed to this bid and the number of items this bid requires, and a cost c, and outputs* 1 *(accept) if and only if $v - k \cdot c \geq t$, where t is a non-negative threshold in $\mathbb{R}$.*

The $0-$threshold composition rule accepts any bidder whose value is greater than the cost to add this bidder to the winning set. For any non-zero threshold t, it means we require a fixed positive welfare in order to accept the bidder.

Basically it means that every unit must individually add positive welfare by accepting this bidder, in order for the bidder to be accepted. Since we send the minimum price out of the k best-priced units available, this ensures that.

3.2 Asymmetric Composition

Here we formally define the core primitive we need, the asymmetric composition. This asymmetric composition will take in a knapsack ranking and a single-unit ranking, along with an asymmetric composition rule, and outputs an allocation rule $x(\cdot, \cdot)$. An asymmetric composition along with a payment rule $p(\cdot, \cdot)$ defines a double auction mechanism. Then, we can prove statements of the form: 'If the ranking algorithms and composition rule satisfy properties A and B, with payment rule C then the double auction mechanism output by their asymmetric composition satisfies property D'. We define an asymmetric composition below.

Definition 35 (Asymmetric Composition) *An asymmetric composition, determined by a knapsack ranking algorithm r, a cost ranking algorithm $\hat{r}$, and a composition rule, greedily determines an allocation as follows. Let n be the number of bidders, t be the number of distributors and m be the total number of units available across all distributors.*

- *For each $j \in [m]$:*
  1. *Let $I_j = \{\}$.*
  2. *For $i \in [n]$: Apply the composition rule on $(r_i(\boldsymbol{b}), \hat{r}_j(\boldsymbol{d}))$. If it outputs* 1 *and $k_i + \sum_{\ell \in I_j} k_\ell \leq j$, $I_j = I_j \bigcup \{i\}$.*
- *For each $j \in [m]$, let $w_j = \sum_{i \in I_j} v_i - \sum_{i \in [j]} \hat{r}_i(\boldsymbol{d})$.*
- *Let j^* be the j with the largest corresponding w_j for $j \in [m]$. Output I_{j^*} and the set of distributors with costs $\{\hat{r}_j(\boldsymbol{d})\}_{j \in [j^*]}$.*

This defines our asymmetric composition given two ranking algorithms. Note that we can also define a slightly different asymmetric composition denoted a *trade-reduced* asymmetric composition. In a trade-reduced asymmetric composition, we remove the least efficient distributor *that was selected by the algorithm.* This distributor will set the price.

Definition 36 (Trade-Reduced Asymmetric Composition) *A Trade-Reduced Asymmetric Composition runs exactly as Definition 35 except after the algorithm we run one additional step:*

- *Remove all units from the least-efficient (worst-ranked) distributor from the accepted set. Also remove any bidders that no longer fit after this distributor is removed.*

When creating a composition using a threshold asymmetric composition rule (Definition 34), we will see that it will be necessary to trade-reduce our composition in order to achieve incentive compatibility.

4 Incentive Compatibility

In this section, we will determine incentive compatibility of knapsack double auctions defined via our composition framework. Specifically, we will show what properties our ranking algorithms and composition rule must follow so that the resulting double auction is truthful for both sides (no party has an incentive to misreport their true value) and also weak group strategy proof (for every group formed, at least one member of the group has no incentive to join it).

4.1 Properties

We start by defining some properties that will be necessary to show both truthfulness and WGSP. We denote a player to be any bidder or distributor participating in the auction.

Definition 41 (Monotone Knapsack Ranking) *We denote a knapsack ranking algorithm to be monotone if a player's ranking both:*

- *(Weakly) Improves with quality: For every bidder i, and bid set $\boldsymbol{b}$, if bidder i changes its bid $b_i = (v_i, k_i)$ to a bid $b_{i'} = (v_{i'}, k_i)$ with $v_{i'} > v_i$, it follows that $r_i(b_{i'}, \boldsymbol{b}_{-i}) \leq r_i(\boldsymbol{b})$ (defined symmetrically for distributors).*
- *(Weakly) Deteriorates with size: For every bidder i and bid set $\boldsymbol{b}$, if bidder i changes its bid $b_i = (v_i, k_i)$ to $b_{i'} = (v_i, k_{i'})$ with $k_{i'} > k_i$, it holds that $r_i(b_{i'}, \boldsymbol{b}_{-i}) \geq r_i(\boldsymbol{b})$ (defined symmetrically for distributors).*

This differs from previous work since we require two properties, both improvement with quality and deterioration with size. This is a natural extension of the monotone ranking algorithm defined in [12], which did not consider single-minded bidders.

Definition 42 (Consistent Knapsack Ranking) *A knapsack ranking algorithm r is consistent if for any two bidders $i \neq i'$ and any bid set $\boldsymbol{b}$, $r_i(\boldsymbol{b}) < r_{i'}(\boldsymbol{b})$ implies that at least one of the following conditions must hold:*

1. $v_i \geq v_{i'}$
2. $k_i \leq k_{i'}$

Here, we must ensure at least one of the two conditions holds for our ranking algorithm to be consistent. Either the value is greater or the size is smaller for one bid to be ranked over another. Notice that bids with different value and different size can be incomparable. In the case where $k = 1$ for all bidders, the definition is equivalent to the consistency definition in [12]. We also apply their definition of monotonicity and consistency for our cost rankings, which depend only on the cost (but inversely) and not on the number of units (since the distributors are not single minded). We do not redefine them here, but monotonicity would require only the first point of Definition 41, and consistency requires only the first point of Definition 42.

Definition 43 (Monotone Asymmetric Composition) *Consider two different inputs to a asymmetric composition rule, with both bidders having same size n: inputs (b_i, c), $(b_{i'}, c')$, and assume that $v_{i'} \geq v_i$ and for each $i \in [n]$, $k_i \geq k'_i$. The composition rule is monotone if for any such two inputs, if the composition rule accepts (b_i, c) then it accepts $(b_{i'}, c')$.*

These two properties will be important to show our incentive compatibility.

4.2 Truthfulness

Here, we show that we can use our definitions from before to show that a composition that follows certain rules will always be truthful.

Theorem 41 (Truthful Composition). *A trade-reduced asymmetric composition that is monotone and composed of a consistent and monotone knapsack ranking algorithm and a consistent and monotone cost ranking algorithm, using a monotone composition rule and applying threshold payments is a truthful double auction mechanism.*

We defer the proof to Appendix C.1.

We cannot protect against sybil attacks unless we include additional assumptions on the distribution. Potentially every distributor would have an incentive to place a bid for one unit of computation for a price immediately less than its own, in order to increase its revenue in the case that it can set its own price. For this work, we do not consider sybil attacks.

4.3 Weak Group Strategy Proofness

Next, we look at how our composition fares against groups of bidders strategically coordinating their bids. Specifically, as we have seen, achieving Group Strategy Proofness is not possible for double auctions [12], therefore, we look to show that our asymmetric composition satisfies weak group strategy proofness.

Theorem 42 (Weak Group Strategy Proofness). *An asymmetric composition of monotone, consistent ranking algorithms with a monotone composition rule and applying threshold payments is a Weak Group Strategy Proof double auction mechanism.*

We defer the proof of this theorem to Appendix C.2.

5 Welfare and Budget Balancing of Compositions

Other than incentive compatibility, there are two more properties of the auction we care about. Welfare and budget balancing. Recall we define these in Section 2. We go into each of these in detail below.

5.1 Welfare

First, we look at the social welfare of the double auction mechanism resulting from our asymmetric composition. Precisely, given the guarantees of welfare achieved by single sided auctions defined by each ranking algorithm, we would like to make a statement about the social welfare of the composition of both into our double auction. For example, in [12] they show that for single unit auctions, a composition of two second price auctions (which provide optimal social welfare) into a double auction provides optimal social welfare. Recall that for knapsack auctions there is no efficient algorithm that solves for the optimal solution in polynomial time, so we will be most likely working with approximations in order to keep the algorithm efficient. In this subsection, we lower bound the welfare of the double auction mechanism resulting from our asymmetric composition below.

To do this, we first must define a set of parameters which will be necessary in doing so.

Quantifying how exhaustive the composition rule is: We will quantify how exhaustive the composition rule by a parameter which tells us how large the set of bidders and distributors is when compared to the size of the set of bidders and distributors in the optimal allocation would be. Specifically, we use $s(\boldsymbol{b}, \boldsymbol{d})$ to output a number between 0 and $\sum_{i\in[n]} k_i$ to be the number of items allocated by our composition. We will denote $s'(\boldsymbol{b}, \boldsymbol{d})$ to denote how many items would have been allocated with the same ranking algorithms and an unconstrained composition rule, and finally we will denote $s^*(\boldsymbol{b}, \boldsymbol{d})$ to be the number of items the optimal solution would have allocated. Concretely, in our case this will quantify how much damage the trade reduction does to our welfare.

Quantifying how close to optimal our ranking algorithms are: We will use parameters $\alpha, \beta \geq 1$ to quantify how close to optimal the solution to the one sided algorithms are for any number q of allowed allocation items. Let us define $v_{\mathsf{OPT}}(q)$ to be the value of the feasible solution with at most q bidders that maximizes total value (analogously, we define $c_{\mathsf{OPT}}(q)$ for distributors, but minimizing cost). Then, we say that a ranking algorithm ALG is an α approximation of the optimal if for any $q > 0$, for any bid vector $\boldsymbol{b}$,

$$v_{\mathsf{ALG}}(q) \geq \frac{1}{\alpha} \cdot v_{\mathsf{OPT}}(q)\,.$$

Similarly for distributors we say that a ranking algorithm is a β approximation of optimal if for any q, for any distributor costs report $\boldsymbol{d}$,

$$c_{\mathsf{ALG}}(q) \leq \beta \cdot c_{\mathsf{OPT}}(q)\,.$$

Note that our ranking algorithm is *perfect* if $\beta = 1$.

Quantifying how difficult the problem instance is: The parameter $\gamma(\boldsymbol{b}, \boldsymbol{d}) = v_{\mathsf{OPT}}(s^*(\boldsymbol{b}, \boldsymbol{d})/c_{\mathsf{OPT}}(s^*(\boldsymbol{b}, \boldsymbol{d}))$ quantifies how hard the problem is by taking the ratio of the optimal cost profile for distributors and optimal value profile for distributors. It has been used also in prior works as the same measurement in related problems [12,25]. A large ratio means that there is a lot more value than cost and therefore there are likely to be a surplus of high value bids and therefore approximating the problem instance is not as hard. It also means that the optimal welfare is large. When there is a ratio close to one the optimal solution is close to 0 and tight, and therefore harder to approximate.

Proving Welfare of an Asymmetric Composition Now we use these to show our theorem on welfare, which uses the parameters above to quantify the optimality of our asymmetric composition.

To do this, first we define a *linear* knapsack ranking algorithm and a *perfect* standard ranking algorithm.

Definition 51 (Linear Knapsack Ranking) *A linear knapsack ranking algorithm is a ranking algorithm which sorts user bids of the form $b_i = (v_i, k_i)$ by using as a sorting key the expression $C(v_i/k_i)$ for any $C \geq 1$.*

This means that the bids with the largest value over size ratio (with some multiplicative weight constant greater than one) are ranked first.

Lemma 51 *A linear knapsack ranking is monotone and consistent.*

Proof. It is easy to verify that our linear knapsack ranking satisfies both our properties.

Definition 52 (Perfect Cost Ranking) *A perfect cost ranking in the unconstrained standard setting is a distributor ranking where for we use only cost to rank the distributors. We sort distributors by smallest cost and output $\hat{r}(\boldsymbol{d})$ such that the $i-th$ element and we rank distributors by smallest cost. We call it perfect since it achieves optimal welfare from the point of view of distributor selection.*

Now we give our theorem. For the theorem below we consider only auction mechanisms which execute at least one trade, meaning $s(\boldsymbol{b}, \boldsymbol{d}) \geq 1$.

Theorem 51 (Approximate Welfare). *Consider input $(\boldsymbol{b}, \boldsymbol{d})$ for which $\mathsf{OPT}(\boldsymbol{b}, \boldsymbol{d}) \geq 0$. The asymmetric composition of a linear knapsack ranking algorithm that is an α approximation of the optimal assignment, and a perfect cost ranking, using a monotone composition rule achieves welfare at least:*

$$\frac{s(\boldsymbol{b}, \boldsymbol{d})}{s'(\boldsymbol{b}, \boldsymbol{d})} \cdot \frac{\frac{\gamma(\boldsymbol{b}, \boldsymbol{d})}{\alpha} - 1}{\gamma(\boldsymbol{b}, \boldsymbol{d}) - 1} \cdot \mathsf{OPT}(\boldsymbol{b}, \boldsymbol{d})\,.$$

We defer the proof of Theorem 51 to Appendix C.4.

The steps in the proof follow along the same lines as the proof in [12], except the arguments are different since we use a different composition rule and different ranking algorithms.

5.2 Budget Balancing

The second property we examine is the budget balancing of the double auction mechanism resulting from our asymmetric composition. Specifically, while we cannot guarantee strong budget balancing (the auctioneer will not gain or lose money from the trade), what we can show is that the double auction mechanism resulting from an asymmetric composition is weak budget balanced (the auctioneer will not lose money from running the auction. We show this below:

Theorem 52 (Weak Budget Balanced Asymmetric Composition). *A trade-reduced monotone asymmetric composition of a linear knapsack ranking and a consistent, monotone standard ranking algorithms using a monotone composition rule and applying threshold payments is weak budget balanced.*

We defer the proof of this theorem to Appendix C.3.

With this, we have shown that our composition can be weak budget balanced when using the correct ranking algorithms. Next, we look at how to use this.

6 Concrete Instantiations and Applications

Now, we move to look at what these ranking algorithms actually look like. It is clear to see that for the single-unit ranking we can always just sort by cost. This ranking is monotone and consistent. [5] Similarly, for our knapsack algorithm, notice that any sorting of users using as sorting value some expression of the form $f(v_i)/g(k_i)$ for any functions f, g would satisfy the consistent and monotone ranking definitions. However, we are only able to show approximate welfare optimality and weak budget balancing for linear knapsack ranking algorithms.

Specifically, if we instantiate our knapsack ranking to use plainly v_i/k_i in order to sort each bid. Ranking items in this fashion for the knapsack problem has been shown to give a $1/2$ approximation to the best knapsack solution for any size, meaning this ranking achieves an $\alpha = 2$ approximation in the parameters defined in Section 5. The sort-by-cost ranking for distributors achieves a $\beta = 1$ optimal solution which then means that a composition of these two algorithms provides a robust auction mechanism with welfare that is close to $1/2$ of the optimal welfare for problem instances with large enough γ. We formalize this below (recall the definition of OPT from Section 2):

Lemma 61 *For any user bid vector* $\boldsymbol{b}$ *and distributor cost bid vector* $\boldsymbol{d}$*, an asymmetric composition by sorting user bids by* v_i/k_i *as our knapsack ranking algorithm and a cost ranking algorithm sorting distributors by their cost* c_j*, using the* $0-$*threshold composition rule and threshold payments is a robust double auction mechanism which achieves welfare at least:*

$$\frac{\frac{\gamma(\boldsymbol{b},\boldsymbol{d})}{2} - 1}{\gamma(\boldsymbol{b},\boldsymbol{d}) - 1} \cdot \mathsf{OPT}(\boldsymbol{b},\boldsymbol{d}) \, .$$

Proof. It follows directly from previous results shown here plus the fact that our knapsack ranking algorithm can approximate knapsack up to a factor of $1/2$.

Notice that for large $\gamma(\boldsymbol{b},\boldsymbol{d})$ this approaches a $1/2$ approximation of the optimal welfare.

One can also define a more general theorem. Any linear knapsack ranking algorithm which provides an α approximation to knapsack, paired with the optimal distributor sorting algorithm, can be converted into a robust double auction mechanism with a $1/\alpha$ approximation to the optimal welfare for problem instances with large enough γ.

6.1 Applications of Interest

Our knapsack double auctions can be useful for a suit of diverse online auctions. In the blockchain space, where there is typically not one provider, but a heterogenous network of providers for computation, storage, among other things, knapsack auctions become very relevant.

[5] We leave it more general in case other constraints cause some distributors to be more appealing than others.

As is, our auctions can cover any setting where the bidders have some knapsack constraint on the amount of resource they want, and distributors don't. As modeled, it is also necessary that bidders don't care whether their service is provided by one or many distributors, as long as it is satisfied. Although most modern services are parallelizable, it could be also very interesting to expand our work to the case where we do not actually want a user's bid to be split among different distributors. Another interesting thing to consider would be whether we can have a truthful mechanism that optimizes only for bidder welfare. Or maybe dropping distributor truthfulness in favor of some notion of fairness for distributors (on how often they get to sell).

One technicality that is necessary to be dealt with is sybil attacks. We assume that the usage of the network is by known clients with a reputation. The auctioneer which is in charge of the auction can work to detect and dissuade sybil attacks through authentication, staking, and slashing.

7 Protocol Modifications and Tradeoffs

In Appendix B, we discuss additional modifications to the protocol which might be of interest in practice. Specifically, we discuss a property orthogonal to some which we discuss in that paper, which is the issue of *fairness*. Fairness, informally, refers to any distributor willing to provide a unit for a price less than the agreed-upon price having equal chance of being assigned it (proportional to the distributor's size). This comes up because for distributors, one important factor for willingly participating in auctions is predictability of the expected return. This is talked about by FileCoin storage providers as paramount for their business [1]. Thus, we believe fairness is a good way to capture this. Unfortunately, fairness is at odds with other desired properties of auctions. Again, we define fairness formally and discuss it in more depth in Appendix B.

8 Conclusion

In this work, we pushed our understanding of double auctions by providing a composition system for knapsack double auctions. Our main motivation for the exploration was a zero-knowledge proof market, although we identified many other settings where such a double auction could be useful. Many questions still remain open in this setting. The two we find more compelling are:

- Can we achieve a closer approximation to the optimal welfare? Ideally we could achieve a $(1 - \epsilon)$ approximation of the welfare for any epsilon. Our framework outlines what is necessary for such result: a monotone and consistent ranking algorithm that can achieve such approximation, along with a way to relate this algorithm to the welfare (if it is not linear).
- Can we provide robust double auctions when both bidders and distributors have a knapsack constraint? And how does this affect the welfare? This setting can be relevant when bidders want their units to come from a single provider (an example could be non-parallelizable computation).

References

1. FileCoin Docs | ROI, July 2023.
2. Gagan Aggarwal and Jason D. Hartline. Knapsack auctions. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm - SODA '06*, pages 1083–1092, Miami, Florida, 2006. ACM Press.
3. Lawrence M Ausubel. An Efficient Ascending-Bid Auction for Multiple Objects. *American Economic Review*, 94(5):1452–1475, December 2004.
4. Lawrence M. Ausubel and Paul R. Milgrom. Package Bidding Vickrey vs. Ascending Auctions. *Revue économique*, 53(3):391–402, 2002. Publisher: Sciences Po University Press.
5. Brian Baisa and Simon Essig Aberg. Quantifying the Inefficiency of Multi-unit Auctions for Normal Goods, 2021.
6. Liad Blumrosen and Noam Nisan. Combinatorial Auctions. In Eva Tardos, Noam Nisan, Tim Roughgarden, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, pages 267–300. Cambridge University Press, Cambridge, 2007.
7. Shuchi Chawla. Lecture Notes in Algorithmic Game Theory, 2011.
8. Yeon-Koo Che and Jinwoo Kim. Optimal collusion-proof auctions. *Journal of Economic Theory*, 144(2):565–603, March 2009.
9. Xianhao Chen, Guangyu Zhu, Haichuan Ding, Lan Zhang, Haixia Zhang, and Yuguang Fang. End-to-End Service Auction: A General Double Auction Mechanism for Edge Computing Services. *IEEE/ACM Trans. Netw.*, 30(6):2616–2629, June 2022.
10. Leon Yang Chu. Truthful Bundle/Multiunit Double Auctions. *Management Science*, 55(7):1184–1198, 2009. Publisher: INFORMS.
11. Edward H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971. Publisher: Springer.
12. Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. Modularity and greed in double auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*, EC '14, pages 241–258, New York, NY, USA, June 2014. Association for Computing Machinery.
13. Theodore Groves. Incentives in Teams. *Econometrica*, 41(4):617–631, 1973. Publisher: [Wiley, Econometric Society].
14. Sujit Gujar and Y. Narahari. Optimal Multi-unit Combinatorial Auctions with Single Minded Bidders. In *2009 IEEE Conference on Commerce and Enterprise Computing*, pages 74–81, July 2009. ISSN: 2378-1971.
15. A-Long Jin, Wei Song, Ping Wang, Dusit Niyato, and Peijian Ju. Auction Mechanisms Toward Efficient Resource Sharing for Cloudlets in Mobile Cloud Computing. *IEEE Transactions on Services Computing*, 9(6):895–909, November 2016. Conference Name: IEEE Transactions on Services Computing.
16. Bernhard Kasberger and Alexander Teytelboym. The Combinatorial Multi-Round Ascending Auction, February 2024. arXiv:2203.11783 [cs, econ].
17. Peyman Khezr and Anne Cumpston. Multi-Unit Auctions: A Survey of Theoretical Literature, May 2020.
18. Kevin Leyton-Brown. Lecture Notes - Multiunit Auctions, 2008.
19. S. Loertscher and C. Mezzetti. A Multi-Unit Dominant Strategy Double Auction. 2013.
20. Simon Loertscher and Claudio Mezzetti. A dominant strategy double clock auction with estimation-based tâtonnement. In *Theoretical Economics*, volume 16, pages 943–978, 2021. ISSN: 1933-6837 Issue: 3 Journal Abbreviation: TE.

21. R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450, April 1992.
22. John McMillan. Selling Spectrum Rights. *The Journal of Economic Perspectives*, 8(3):145–162, 1994. Publisher: American Economic Association.
23. Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, April 1983.
24. Jer Shyuan Ng, Wei Yang Bryan Lim, Zehui Xiong, Dusit Niyato, Cyril Leung, and Chunyan Miao. A Double Auction Mechanism for Resource Allocation in Coded Vehicular Edge Computing. *IEEE Transactions on Vehicular Technology*, 71(2):1832–1845, February 2022. Conference Name: IEEE Transactions on Vehicular Technology.
25. Tim Roughgarden and Mukund Sundararajan. Quantifying inefficiency in cost-sharing mechanisms. *J. ACM*, 56(4):23:1–23:33, July 2009.
26. William Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, 16(1):8–37, 1961. Publisher: [American Finance Association, Wiley].
27. Wenhao Wang, Lulu Zhou, Aviv Yaish, Fan Zhang, Ben Fisch, and Benjamin Livshits. Mechanism Design for ZK-Rollup Prover Markets, April 2024. arXiv:2404.06495 [cs].

A Related Work

Auctions have been thoroughly studied in many different settings. In the interest of clarity and completeness, we outline below some settings we found to be similar to the prover network setting we study here, although with some important differences which we also discuss.

A.1 Combinatorial/Knapsack Auctions

The study of combinatorial auctions was initially studied by McMillan [22] in the setting where there are several different item types and users might be interested in only a set of them (these are also called package auctions). This setting was also studied in other works [4,6,16]. Along with the fact that this setting is not exactly equivalent to ours where users are interested in varying quantities of 'the same' item, these works also only talk about one-sided auctions. Two works were identified that work in the setting of multiple sellers. The work by Gujar and Narahari [14] and the work by Chu [10]. In [14] the mechanism is not budget balanced and in [10] to the best of our understanding the welfare analysis does not extend to the case where sellers sell multiple items of the same kind, which is the case in our setting.

A.2 Multi-Unit Auctions for Bidders with Diminishing Marginal Value per Item

Another setting that has been extensively studied is the setting where bidders have an interest in many items, but have marginal decreasing value for each (if they have a value of v_1 for acquiring the first item, they have value $v_2 < v_1$ for the second one and so on). In this setting, the VCG auction fails in both budget balancing and computational efficiency, the optimal solution is intractable to compute. Many clever solutions are found to balance the properties and achieve something that 'almost' achieves all properties (given that achieving all of them is impossible [23]) [3,17,5]. One work has also studied this problem in the double auction setting [19,20] but the arguments all require the property of diminishing marginal value per item.

A.3 Concurrent Related Work in Auctions for Prover Networks

Wang et al. [27] considered the setting of a prover network and also examine some mechanisms in this scenario. Their work, however, focuses on a profit-maximizing middleman (in our work, called a gateway), which in turn sacrifices some other properties deemed desired by previous works. As we will see, it is extremely hard to conciliate a profit-maximizing gateway with the set of desiderata we define for the prover network. In communication with the authors they have made considerable changes to the model they work on in a new unreleased version of their work, and while related the models are different.

B Fairness and Disincentivizing Overbidding by Sellers

We have shown that a composition framework that allows us to input two ranking algorithms with certain properties to output a double auction mechanism that is truthful for both bidders and distributors; in other words, the profit maximizing strategy for both operators and distributors is to report their private values and capacities to the auctioneer. Distributors are not incentivized to bid more items than they have because they will not receive any reward for selling items they cannot provide. However, we run into a problem where, although bidding truthfully is a Bayesian Nash Equilibrium, overbidding is also a Bayesian Nash Equilibrium, since it also maximizes revenue (the distributor doesn't actually lose anything it didn't already not have by overbidding).

Another issue we might consider is the case where many distributors have the same cost bid, or many distributors are willing perform computation for the price agreed upon by the auction (even if they were not the best option in terms of welfare), it could be of interest to reduce the sum of the welfare of the operators in exchange for fairness, where fairness would refer to how likely an operator is to be picked given that its price was below the price that is being paid by the bidders. We will call this property fairness. We will say that an auction mechanism is fair if any distributor whose cost per unit is lower than the agreed upon auction price has a chance of being picked proportional to its size.

In this section, we first define fairness formally. Then we discuss how it is at odds with welfare and truthfulness in some regards, but is interesting to consider for real world applications. We then show a way to devise a fair scheme that also disincentives overbidding (at the cost of some welfare).

We consider fairness on a per-unit basis, ignoring which distributor the bid came from.

Definition B1 (Fairness) *We denote an auction mechanism to be fair if for any two sellers with costs c_i, c_j respectively; if c_i and c_j are smaller than the cut-off sell price c, then the probability that seller i is assigned is the same probability that seller j is assigned.*

Consider any scenario where the demand is much higher than the supply for a certain item. In this case, competing bidders will drive up the price well above the cost each operator has for a given item. When this happens, there can be an optimal assignment where different sets of distributors are valid assignments. In the current setting, we always pick the most efficient distributors (distributors with smallest cost) to participate. However, this is not necessarily fair. Although this is welfare maximizing for the auction round, willing bidders cannot participate because other distributors are more efficient (given the knapsack constraint, in high demand there may be sellers who are willing to provide units but are left out).

In this setting, we can see that our current mechanism is not *fair* as defined in Definition B1. Although it is in detriment to total welfare to pick a seller j over seller i if $c_j > c_i$ (since it is subtracted from the welfare), fairness is also an important property in settings where sellers (participants in a recurring auction mechanism for example) need a predictable and stable allocation to be able to keep providing units in the long run [1]. Fairness can help reduce the variance in seller revenue by ensuring that they are allocated with probability proportional to the amount of compute they are committed to provide, whenever cost allows.

The question then is whether we can modify our composition to account for fairness, and how will that impact welfare? To analyze this, we propose a different notion of welfare, namely, to consider only bidder welfare, and consider seller fairness rather than seller welfare. This way we can still argue about some welfare guarantee for bidders when talking about fair mechanisms, which by definition will not maximize welfare as defined in Section 2. Notice that we are only modifying the welfare property and adding the fairness property, but would still like the mechanism to satisfy our incentive compatibility properties as well as budget balancing.

Definition B2 (Random Asymmetric Composition) *Let us define a fair asymmetric composition to work exactly as in Definition 36 except after the trade reduction we don't pick the most efficient distributors to participate and instead sample uniformly each participant out of the costs per unit c_j such that $c_j < c$ where c is the price determined by the threshold payment rule.*

Notice that by definition, this randomized asymmetric composition is fair. However, it *does not* maintain the properties we had shown before from our standard asymmetric composition. Specifically, it is easy two see two thing immediately:

– The auction is no longer welfare maximizing, since the participating distributors are no longer the distributors with the least cost.
– The auction is no longer truthful, since now efficient distributors are incentivized to overbid the amount of compute they can provide. Notice that now distributors that were fully allocated in the previous composition function by being the most efficient, might now only be partly allocated. This means they can overbid the amount they have to increase their chances of being picked by the randomized algorithm and have all the units they have available allocated (by bidding more than what they have). Clearly this is not a truthful auction.

Both of these problems can be addressed by introducing staking and slashing into the distributors' reward function. Specifically, we require each distributor d_j to deposit some value $h_j \cdot u_j$ before the auction. Then, whenever a distributors cannot provide the entire amount of units of compute they bid, this amount can be taken by the auctioneer and burned. (If this does not happen, the distributor gets its deposit back at the end of the auction.) Now, we can adjust h_j to be such that the expected welfare of the distributors is maximized when bidding truthfully. Although it is not straightforward to compute the optimal staking and slashing mechanism such that this is true, it is easy to see that we can always find some mechanism for which this is true by making the stake large enough. We leave a further development of fairness and a more formal analysis of it to future work.

C Deferred Proofs

In this section we provide the proofs deferred from the main body.

C.1 Proof of Theorem 41

We first restate the theorem then provide the proof.

Theorem 41 (Truthful Composition). *A trade-reduced asymmetric composition that is monotone and composed of a consistent and monotone knapsack ranking algorithm and a consistent and monotone cost ranking algorithm, using a monotone composition rule and applying threshold payments is a truthful double auction mechanism.*

Proof. We will consider truthfulness for bidders and distributors.

Truthful for bidders: First, notice that any participating bidder has nonnegative utility. Winning bids have utility greater than or equal to 0, and losing bids have utility 0. Consider any bidder i with $b_i = (v_i, k_i)$ who bid his truthful values. Let us consider two cases:

1. **Bidder i was accepted:** In this case, the bidder i has utility greater than or equal to 0. Notice i has no incentive to reduce its value v_i since the

threshold payment ensures that the payment by i is the smallest it had to bid to get accepted. Furthermore, reducing k_i is of no benefit to bidder i since receiving less than k_i items has value 0 for this bidder. Increasing the bids cannot increase its utility either.

2. **Bidder i was not accepted:** In this case, i has utility 0. Notice it has no incentive to change its reported v_i, since by definition the threshold payment necessary to be accepted for i was *at least* v_i (to get accepted it would get at most utility 0), and by consistency, for every i' ranked higher than i, either $v_i \geq v_{i'}$ or $k_i \leq n_{i'}$. Furthermore, notice that by monotonicity of the ranking increasing the k_i (or decreasing v_i) can not give bidder i a better ranking. Finally, decreasing k_i brings the value of the package to 0 value to i so this cannot increase utility.

Truthful for distributors: Again, notice that any participating distributor has non-negative utility. They can either be picked and get utility greater than or equal to 0, or not be picked and get utility 0. Consider any distributor with bid $d_i = (u_i, c_i)$, its number of units provided and cost per unit, respectively. We again consider two cases:

1. **Distributor i was accepted:** By threshold payments, we note that distributor i has no incentive to increase its reported cost, since it gets paid proportional to the maximum it could have bid and still be accepted.
2. **Distributor i was not accepted:** By monotonicity of the ranking, distributor i cannot improve its ranking by reporting a higher cost. Furthermore, by the threshold payment rule the cost received by accepted distributors was at most c_i, which means that even if this distributor bid some $c' < c_i$ (where c_i is its true cost) to participate, it could achieve at most utility 0 (which it already has).

Notice that this only holds because we remove the least efficient distributor, or else a distributor on the border could potentially set its own price by being both an accepted bid and a bid that sets the price to receive (only removing a distributor if it has both accepted and not accepted units also does not work, since then this distributor has an incentive to reduce the number of units it bids to not go over the assignment). □

C.2 Proof of Theorem 42

We first restate the theorem for WGSP compositions, then provide the proof.

Theorem 42 (Weak Group Strategy Proofness). *An asymmetric composition of monotone, consistent ranking algorithms with a monotone composition rule and applying threshold payments is a Weak Group Strategy Proof double auction mechanism.*

Proof. First, we show that for any coalition composed of multiple bidders and multiple distributors, at least one participant is indifferent to participating in the coalition.

Multi-bidder coalitions: We can split bidder coalitions into coalitions of losing bidders, coalitions of winning bidders, and a coalition with a mix of both.

A coalition of losing bidders has utility 0. Any strategy that involves some player lowering their bid would make this player's utility the same as without colluding (0) and therefore this player would always be indifferent to participating. Likewise, any strategy which involves some player increasing their bid in order to get accepted would make this player have either 0 or negative utility, meaning this player is indifferent to joining the coalition. Finally, any bid which does not change the accepted set does not change the coalition's utility.

A coalition of winning bidders has utility greater than or equal to zero. Notice that any strategy that involves some party lowering bids to anything greater than the payment threshold would keep this party's payment (and therefore utility) unchanged and so this party would be indifferent to participating. Furthermore, any strategy that involves some party lowering their bid until they get rejected would make this party's utility 0, which is less than or equal to its utility as a winning bidder, therefore this party is indifferent to participating. Increasing a bid cannot increase winning bidder's utility by the threshold payment rule.

In a coalition of both winning and losing bidders, notice that no winning bidder can lower their bid further than the threshold (which would make its utility 0 and therefore less than or equal to its current utlity). However, to achieve non-zero utility, a losing bidder must be accepted, but this can only happen if a member of the coalition that is a winning bidder plays a losing bid, which as we just argued would make this member's utility less than or equal to his utility when not participating. Therefore, in a coalition of both winning and losing bidders, at least one losing bidder or one winning bidder is indifferent (or worse) to participating.

Multi-distributor coalitions: The argument for multi-distributor coalitions follows almost exactly symmetrically to the argument for bidder-only coalitions, so we omit it.

Mixed coalitions: Now, we move to show that in any coalition with both bidders and distributors. Notice that any coalition with more than one distributor or more than one bidder is already shown to have at least one indifferent participant by the arguments above. Therefore, what remains is to show that for any coalition of a single bidder and a single distributor, at least one participant is indifferent to participating in the coalition.

The argument follows similarly to our argument for coalitions of winning and losing bidders:

- Coalition of a bidder with a winning distributor: The distributor cannot increase its utility by bidding a smaller cost (will receive the same payment by the threshold rule). By bidding a larger cost so that it is no longer a winning distributor, this distributor's utility will be at most the same as before. Then, the only strategy the distributor would not be indifferent to is bidding a larger cost and still be accepted. However, the larger cost that the distributor can bid is at most the threshold payment and with that bid it would have its utility unchanged. So for any coalition of a single bidder with a single distributor, this distributor is indifferent.

– Coalition of a bidder with a losing distributor: follows as above.

Then, we have shown that for any coalition, at least one party is indifferent to participating in the coalition, and this concludes our proof for Weak Group Strategy Proofness. Again, this holds because there is no distributor or bidder able to set the price and participate. □

C.3 Proof of Theorem 52

Again, we restate our theorem for weak budget balancing and then provide the proof.

Theorem 52 (Weak Budget Balanced Asymmetric Composition). *A trade-reduced monotone asymmetric composition of a linear knapsack ranking and a consistent, monotone standard ranking algorithms using a monotone composition rule and applying threshold payments is weak budget balanced.*

Proof. We apply per-unit threshold payment, where each bidder pays the minimum per-unit price that was bid and rejected.

Notice that when we trade-reduce our composition, we remove one active distributor, and all bidders that were assigned to that distributor for at least one unit. This means we remove at least one bidder. Let $b_\ell = (v_\ell, k_\ell)$ be the lowest ranked bidder that was removed by the trade reduction, and let c' be the per unit cost of the distributor that was removed by the trade reduction. Notice that by our asymmetric composition, we know that $v_\ell/k_\ell \geq c'$. We also know that for any per-unit cost of distributors still active after the trade reduction, their cost $c_a \leq c'$ (this holds because our distributor ranking algorithm is monotone and consistent). Likewise, for any accepted bidder, we have that for their bid $b_a = (v_a, k_a)$, it holds that $v_a/k_a \geq v_\ell/k_\ell$. This is by definition of our linear knapsack ranking algorithm.

Now, notice that for any accepted bid b_a, if the bidder reduces its bid such that $v'_a/k_a \leq v_\ell/k_\ell$, it will be ranked after the losing bid, and therefore be assigned the highest cost distributor and be trade-reduced (no longer accepted). This holds by definition of our ranking algorithm and monotonicity of our composition rule. This means that each accepted bidder will pay *at least* v_ℓ/k_ℓ per unit. This is by definition of threshold payments. Similarly, by the same argument we see that each distributor will get paid *at most* c' per unit . Since $v_\ell/k_\ell \geq c'$, it follows that our composition is weak budget balanced. □

C.4 Proof of Theorem 51

Finally, we restate the theorem and provide the proof for our welfare theorem. We use the definition of OPT from Section 2.

Theorem 51 (Approximate Welfare). *Consider input $(\boldsymbol{b}, \boldsymbol{d})$ for which $\mathit{OPT}(\boldsymbol{b}, \boldsymbol{d}) \geq 0$. The asymmetric composition of a linear knapsack ranking algorithm that is an α approximation of the optimal assignment, and a perfect cost ranking, using a monotone composition rule achieves welfare at least:*

$$\frac{s(\boldsymbol{b}, \boldsymbol{d})}{s'(\boldsymbol{b}, \boldsymbol{d})} \cdot \frac{\frac{\gamma(\boldsymbol{b}, \boldsymbol{d})}{\alpha} - 1}{\gamma(\boldsymbol{b}, \boldsymbol{d}) - 1} \cdot \mathit{OPT}(\boldsymbol{b}, \boldsymbol{d}) \,.$$

Proof. The final goal is to show that $v_{\mathsf{ALG}}(s(\boldsymbol{b}, \boldsymbol{d})) - c_{\mathsf{ALG}}(s(\boldsymbol{b}, \boldsymbol{d})) \geq \frac{s(\boldsymbol{b}, \boldsymbol{d})}{s'(\boldsymbol{b}, \boldsymbol{d})} \cdot \frac{\frac{\gamma(\boldsymbol{b}, \boldsymbol{d})}{\alpha} - \beta}{\gamma(\boldsymbol{b}, \boldsymbol{d}) - 1} \cdot (v_{\mathsf{OPT}}(s^*(\boldsymbol{b}, \boldsymbol{d})) - c_{\mathsf{OPT}}(s^*(\boldsymbol{b}, \boldsymbol{d})))$.

Now recall $s(\boldsymbol{b}, \boldsymbol{d})$ represents the total number of units transacted. This is the sum of all the sizes of the bidders accepted, $\sum_{i \in B} k_i$. Now, let us split each v_i accepted into k_i values $v_1^i, \ldots, v_{k_i}^i$ where each $v_\ell^i = v_i / k_i$, for $\ell \in [k_i]$. We do this for each accepted bid (v_i, k_i) and concatenate the outputs into a single vector $\boldsymbol{w}$ of size $\sum_{i \in B} k_i$. Let $m = \sum_{j \in |\boldsymbol{d}|} u_j$. For conciseness, let us also use the z to represent $\hat{r}_1(d), \ldots, \hat{r}_m(d)$, the costs for each unit reported by the distributors sorted inversely, such that $z_1 \leq z_2 \leq \ldots \leq z_m$.

Now, we can write the welfare of our output as follows:

$$v_{\mathsf{ALG}}(s(\boldsymbol{b}, \boldsymbol{d})) - c_{\mathsf{ALG}}(s(\boldsymbol{b}, \boldsymbol{d})) = \sum_{i=1}^{s(\boldsymbol{b}, \boldsymbol{d})} w_i - z_i \,,$$

and the same holds when substituting $s(\boldsymbol{b}, \boldsymbol{d})$ for $s'(\boldsymbol{b}, \boldsymbol{d})$.

$$\begin{aligned}
v_{\mathsf{ALG}}(s(\boldsymbol{b}, \boldsymbol{d})) - c_{\mathsf{ALG}}(s(\boldsymbol{b}, \boldsymbol{d})) &= \sum_{i=1}^{s'(\boldsymbol{b}, \boldsymbol{d})} (w_i - z_i) - \sum_{i=s(\boldsymbol{b}, \boldsymbol{d})+1}^{s'(\boldsymbol{b}, \boldsymbol{d})} (w_i - z_i) \\
&\geq \sum_{i=1}^{s'(\boldsymbol{b}, \boldsymbol{d})} (w_i - z_i) - (s'(\boldsymbol{b}, \boldsymbol{d}) - s(\boldsymbol{b}, \boldsymbol{d})) \frac{1}{s(\boldsymbol{b}, \boldsymbol{d})} \sum_{i=1}^{s(\boldsymbol{b}, \boldsymbol{d})} (w_i - z_i) \\
&= (v_{\mathsf{ALG}}(s'(\boldsymbol{b}, \boldsymbol{d})) - c_{\mathsf{ALG}}(s'(\boldsymbol{b}, \boldsymbol{d}))) - \left(\frac{s'(\boldsymbol{b}, \boldsymbol{d})}{s(\boldsymbol{b}, \boldsymbol{d})} - 1 \right) (v_{\mathsf{ALG}}(s(\boldsymbol{b}, \boldsymbol{d})) - c_{\mathsf{ALG}}(s(\boldsymbol{b}, \boldsymbol{d}))) \,.
\end{aligned}$$

Here, in the first line we simply split the sum. The second line is smaller than or equal to the first because, by construction of our auction, the sum of the value of the last $s' - s$ bids are smaller than or equal to the $(s' - s)/s$ fraction of the sum first s bids matched (by Lemma 51), and from the second to third line we simply rearrange. Now, with some more rearranging we get the following:

$$v_{\mathsf{ALG}}(s(\boldsymbol{b}, \boldsymbol{d})) - c_{\mathsf{ALG}}(s(\boldsymbol{b}, \boldsymbol{d})) \geq \left(\frac{s(\boldsymbol{b}, \boldsymbol{d})}{s'(\boldsymbol{b}, \boldsymbol{d})} \right) (v_{\mathsf{ALG}}(s'(\boldsymbol{b}, \boldsymbol{d})) - c_{\mathsf{ALG}}(s'(\boldsymbol{b}, \boldsymbol{d}))) \,.$$

Then, also notice that we iterate through every size from in $[m]$, which means that at some point our algorithm considers the optimal size $s^*(\boldsymbol{b}, \boldsymbol{d})$. Then, the welfare on the set output by our vanilla composition (no trade reduction) is at least as large as the welfare of running our estimation on the optimal size. This means that:

$$v_{\mathsf{ALG}}(s'(\boldsymbol{b},\boldsymbol{d})) - c_{\mathsf{ALG}}(s'(\boldsymbol{b},\boldsymbol{d})) \geq v_{\mathsf{ALG}}(s^*(\boldsymbol{b},\boldsymbol{d})) - c_{\mathsf{ALG}}(s^*(\boldsymbol{b},\boldsymbol{d}))\,.$$

Finally, since our linear knapsack ranking is a α approximation of the optimal, and our distributor ranking is a perfect, it follows that:

$$\begin{aligned} v_{\mathsf{ALG}}(s^*(\boldsymbol{b},\boldsymbol{d})) - c_{\mathsf{ALG}}(s^*(\boldsymbol{b},\boldsymbol{d})) &\geq \frac{1}{\alpha} v_{\mathsf{OPT}}(s^*(\boldsymbol{b},\boldsymbol{d})) - c_{\mathsf{OPT}}(s^*(\boldsymbol{b},\boldsymbol{d})) \\ &= \left(\frac{\gamma(\boldsymbol{b},\boldsymbol{d})}{\alpha} - 1\right) c_{\mathsf{OPT}}(s^*(\boldsymbol{b},\boldsymbol{d})) \\ &= \frac{\frac{\gamma(\boldsymbol{b},\boldsymbol{d})}{\alpha} - 1}{\gamma(\boldsymbol{b},\boldsymbol{d}) - 1} \left(v_{\mathsf{OPT}}(s^*(\boldsymbol{b},\boldsymbol{d})) - c_{\mathsf{OPT}}(s^*(\boldsymbol{b},\boldsymbol{d}))\right). \end{aligned}$$

Putting all the inequalities together we prove the initial statement. □