Modularity and Greed in Double Auctions

PAUL DÜTTING, Stanford University
TIM ROUGHGARDEN, Stanford University
INBAL TALGAM-COHEN, Stanford University

Designing double auctions is a complex problem, especially when there are restrictions on the sets of buyers and sellers that may trade with one another. The goal of this paper is to develop "black-box reductions" from double-auction design to the exhaustively-studied problem of designing single-sided mechanisms.

We consider several desirable properties of a double auction: feasibility, dominant-strategy incentive-compability, the still stronger incentive constraints offered by a deferred-acceptance implementation, exact and approximate welfare maximization, and budget-balance. For each of these properties, we identify sufficient conditions on the two one-sided mechanisms — one for the buyers, one for the sellers — and on the method of composition, that guarantee the desired property of the double auction.

Our framework also offers new insights into classic double-auction designs, such as the VCG and McAfee auctions with unit-demand buyers and unit-supply sellers.

Categories and Subject Descriptors: F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

Additional Key Words and Phrases: Trade Reduction Mechanism, Deferred-Acceptance Auctions

1. INTRODUCTION

Double auctions play an important role in mechanism design theory and practice. They are of theoretical importance because they solve the fundamental problem of how to organize trade between a set of buyers and a set of sellers, when both the buyers and the sellers act strategically. Important practical applications include the New York Stock Exchange (NYSE), where buyers and sellers trade shares, and the upcoming spectrum auction conducted by the US Federal Communication Commission (FCC), which aims at reallocating spectrum licences from TV broadcasters to mobile communication providers. The idea of increasing economic efficiency by reallocating items from lower-valued uses to higher-valued ones is reflected in both of these examples and is common to many other double auction applications.

Designing double auctions can be a complex task, with several competing objectives. Even in simple double auction settings with unit-demand buyers and unit-supply sellers it is generally impossible to achieve economic efficiency and dominant strategy incentive compatibility (DSIC) without running a deficit [Hurwicz 1972]. Myerson and Satterthwaite [1983] extended this impossibility result to also include Bayes-Nash incentive compatible (BIC) mechanisms. Since then much of the literature on double auctions has focused on trading-off efficiency, incentive compatibility and budget balance (BB) [e.g., McAfee 1992; Satterthwaite and Williams 1989; Rustichini et al. 1994; Satterthwaite and Williams 2002].

Authors' addresses: P. Dütting, T. Roughgarden and I. Talgam-Cohen, Department of Computer Science, Stanford University, 353 Serra Mall, Stanford, CA 94305, USA; email: {duetting, tim, italgam}@cs.stanford.edu.

EC'14, June 8–12, 2014, Stanford, CA, USA. Copyright © 2014 ACM 978-1-4503-2565-3/14/06 ...$15.00.
http://dx.doi.org/10.1145/2600057.2602854

The complexity of double auction design grows when there are restrictions on the sets of buyers and sellers that may trade with one another. Such restrictions can arise from resource, policy or legal constraints. For example, a requirement for diversity among the allocated buyers translates to a *matroid* constraint. Another example is sellers whose production consumes a limited resource, which leads to a *knapsack* constraint. While the optimization problem associated with the former can be solved optimally in polynomial time, the latter leads to an optimization problem that can only be solved approximately in polynomial time.

Perhaps the cleanest approach to designing double auctions for complex settings is via a "black-box reduction" to the exhaustively-studied problem of designing single-sided mechanisms. The goal of this paper is to develop the theory that explains when and how such black-box reductions work.

We consider several desirable properties of a double auction: feasibility, DSIC, the still stronger incentive constraints offered by a deferred-acceptance implementation such as weak group-strategyproofness (WGSP) [Milgrom and Segal 2014], exact and approximate welfare maximization, and budget balance. For each of these properties, we identify sufficient conditions on the two one-sided mechanisms — one for the buyers, one for the sellers — and on the method of composition, that guarantee the desired property of the double auction.

1.1. Our Approach—Double Auctions via Composition

Our modular approach applies to double auction settings with identical goods, unit-demand buyers and unit-supply sellers, and a wide variety of constraints on which buyers/sellers can be accepted for trade. It decomposes the design task into the tasks of designing algorithms for each side of the market and a composition rule. All our results are for greedy algorithms. The greedy algorithms each return a stream of players, and the composition rule repeatedly queries these algorithms for the next buyer-seller pair, and decides whether this pair should trade or not based on the bids/asks it has seen so far. We distinguish between a forward and a backward way of implementing this: in the forward way we irrevocably accept pairs, while in the backward way we irrevocably reject pairs. In both cases we apply threshold payments.

This decomposition allows us to derive the properties of the resulting double auction from the properties of the greedy algorithms and the composition rule. We establish several composition results:

— *Welfare composition theorem:* This result relates the approximation guarantee of the double auction to the approximation guarantees of the greedy algorithms and the properties of the composition rule. We require that the approximation guarantees are uniform, i.e., they hold for every cardinality, and that the greedy algorithms are consistent, i.e., the players appear in the streams in order of their quality.
— *DSIC composition theorem:* A double auction that is based on composition is DSIC if the greedy algorithms are rank monotone and consistent and the composition rule is monotone. Rank monotonicity means that by improving their report, players can only improve their position in the stream. Monotonicity means that if a pair is accepted and they only improve their reports, they remain accepted.
— *WGSP composition theorem:* This result shows that if both greedy algorithms are deferred-acceptance algorithms and the composition rule bases its decision only on already rejected pairs — e.g., by accepting a pair only if it has already rejected a pair with positive gain from trade — then their backward composition can again be viewed as a deferred-acceptance algorithm and thus inherits the strong incentive properties established in [Milgrom and Segal 2014].

— *BB composition theorem:* According to this result the backward composition of two deferred-acceptance algorithms that only bases its decisions on previously rejected pairs satisfies budget balance.

Our framework yields a simple and transparent method for designing DSIC double auctions, and the first methodology known to date for designing double auctions with incentive properties beyond DSIC. Despite its simplicity it is general enough to capture the two landmark DSIC mechanisms — the Vickrey-Clarke-Groves (VCG) mechanism and McAfee's trade reduction mechanism — and it leads to previously unobserved properties of these mechanisms.

1.2. A New Perspective on Well-Known Mechanisms

Even in the simplest case to which our framework applies — no constraints on which buyers and sellers can be accepted for trade — it casts a new light on well-studied double auction mechanisms:

— *VCG mechanism via composition:* Sort the buyers by value from high to low and the sellers by cost from low to high. Go through the corresponding lists, and accept the next buyer-seller pair if its gain from trade is positive. If we proceed like this we will accept all pairs with positive gain from trade.

— *McAfee's trade reduction mechanism via composition* Sort the buyers by value from low to high and the sellers by cost from high to low. Go through the corresponding list, and reject the next pair if among the already rejected pairs no pair has positive gain from trade. This will lead to the rejection of only one pair with positive gain from trade, namely the one with the lowest gain.

Our composition theorems now imply that VCG is DSIC and efficient, while trade reduction is WGSP, BB and achieves welfare within $1 - 1/t^*$ of the optimal welfare, where t^* is the number of trades in the optimal solution. To the best of our knowledge the fact that trade reduction is WGSP has not been observed before. Another remarkable consequence of our analysis is that trade reduction not only yields the best-possible welfare approximation (in terms of t^*) subject to DSIC and BB, but also subject to the sole requirement of WGSP.

1.3. Applying the Framework — Matroids and Beyond

Apart from these novel implications for already known mechanisms, our framework is applicable to a wide range of feasibility-constrained environments. We assume that the constraint across markets is that there are as many buyers as sellers, but allow different feasibility constraints on either side of the market. For example:

— *Matroid constraints:* A concrete example are situations where policy considerations require that the winning set of buyers in a double auction will not be dominated by buyers of a certain type, as opposed to a more diversified group of buyers. Imbalance of this form can be bounded by imposing a ceiling on the number of winning players of each type, as in a partition matroid.

— *Knapsack constraints:* As an example consider the case where sellers are firms that produce an identical good, but emit different levels of pollution, and there is a cap on the overall pollution placed by the government.

— *Matching constraints:* This constraint arises, for example, when the seller side consists of pairs of firms producing complimentary goods — for instance, almond butter and quince jelly or wheels and trucks — and the buyer side wishes to buy the compound product only. The winning seller pairs must then correspond to edges that form a matching in the bipartite graph with firms as nodes, and edges indicating willingness to work together.

Intuitively, the first setting is precisely the setting in which the greedy algorithm that sorts elements by weight works optimally. The second setting can be thought of as a relaxation of the matroid constraint, in which greedy by weight is not optimal but often performs well. What both settings have in common is that the greedy algorithm can also be implemented as a deferred-acceptance algorithm. In the third setting greedy by weight is still near optimal, but it is not clear how to implement it as a deferred-acceptance algorithm. We show that a slightly more complicated deferred-acceptance algorithm achieves the same approximation guarantee.

Our framework yields novel VCG and trade reduction-style mechanisms for all three settings that are either DSIC or WGSP, respectively. It also translates the approximation guarantees for the greedy algorithms into welfare guarantees for the double auction. These guarantees show that the welfare degrades gracefully as we move away from settings in which greedy is optimal.

1.4. Related Work

The design principle of modularity is embraced in a diverse range of complex design tasks, from mechanical systems through software design to building architecture [Baldwin and Clark 2000]. Splitting a complex design task into parts or modules, addressing each module separately and combining the modules into a system helps make the design and analysis tractable and robust. Economic mechanisms that operate in complex incentive landscapes while balancing multiple objectives are natural candidates for reaping the benefits of modularity. A notable predecessor of our work in this regard is work by Mu'alem and Nisan [2008], which applies a modular design approach in the context of one-sided combinatorial auctions. They develop an array of truthful algorithmic techniques, and demonstrate how different combinations of these techniques are flexible enough to implement mechanisms for several cases of combinatorial auctions with single-minded bidders.

A first line of related work on double auctions seeks to escape the impossibility results of Hurwicz [1972] and Myerson and Satterthwaite [1983], which show that no double auction mechanism can simultaneously be efficient, DSIC or BIC, and BB by relaxing the efficiency requirement. This direction can be divided into mechanisms that are BIC and mechanisms that are DSIC. An important example of the former is the buyer's bid double auction of Satterthwaite and Williams [1989]; Rustichini et al. [1994]; Satterthwaite and Williams [2002], which sets a single price so as to equal supply and demand. More recent work that falls into this category is [Cripps and Swinkels 2006; Fudenberg et al. 2007]. A prominent example of the latter is McAfee's trade reduction mechanism, which disallows all but the least efficient trade. This mechanism has been generalized to more complex settings in [Babaioff and Nisan 2004; Bredin and Parkes 2005; Gonen et al. 2007; Babaioff et al. 2009; Chu 2009]. More recent work that falls into this category is [Kojima and Yamashita 2013; Blumrosen and Dobzinski 2014] (the work of [Kojima and Yamashita 2013] applies to ex post incentive compatibility as it addresses interdependent values). By Satterthwaite and Williams [2002], the basic variant of the buyer's bid double auction and the trade reduction mechanism are optimal subject to BIC/DSIC and BB. Our work is orthogonal to these works as it focuses on stronger incentives rather than BB, and complements this line of work with the finding that trade reduction is optimal subject to WGSP.

The second line of related work on double auctions is initiated by Carroll [2013], who analyzes the tradeoff between incentives and efficiency while insisting on budget balance. In his work he quantifies the degree to which a mechanism incentivizes truthful reporting in terms of the ex post regret of a player, and relates it to the ex post welfare loss of the mechanism — both in the worst case over all inputs. By weakening the incentive properties, his work moves in the opposite direction than this work.

Our work also builds upon previous work on incentives in single-parameter mechanism design problems. The classic result in this context is due to Myerson [1981] who characterizes DSIC mechanisms as mechanisms that are monotone and apply threshold payments. Although recent progress towards characterizing WGSP mechanisms has been made [Juarez 2013], a similarly simple characterization for WGSP mechanisms is not available. Deferred-acceptance algorithms on which part of our work is based are proposed in [Milgrom and Segal 2014], and their performance is analyzed in [Dütting et al. 2014]. Our work extends the deferred-acceptance framework from one-sided settings to two-sided settings.

1.5. Paper Organization

Section 2 covers preliminaries of the settings to which our analysis applies, and formally defines "output" properties of the double auction mechanisms we are interested in, including incentive compatibility of different degrees (DSIC and WGSP), welfare and BB; this section may be skipped by the expert reader. In Section 3 we describe our composition framework: First we define the corresponding "input" properties of the one-sided modules that are needed to guarantee the required output ones upon composition; then we turn to properties of the composition itself.

The following sections are roughly organized by output property: Section 4 proves the DSIC and WGSP composition theorems by analysis of the one-sided and composition properties that lead to these incentive guarantees; Section 5 has a similar analysis leading to the welfare composition theorem, whose tightness is studied in Section 6. Finally, in Section 7 we prove our BB composition theorem, as well as corresponding lower bounds on the welfare achievable subject to DSIC and BB.

2. PROBLEM STATEMENT

In this section we formally define the double auction settings and the properties of double auction mechanisms that we are interested in. We also present three settings that will serve as running examples.

2.1. Double Auction Settings

We study single-parameter double auction settings: These are two-sided markets, with n buyers on one side of the market and m sellers on the other. There is a single kind of item for sale. The buyers each want to acquire a single unit of this item, and the sellers each have a single unit to sell. Each buyer i has a value $v_i \geq 0$, and each seller j has a cost $c_j \geq 0$. We denote by $\vec{v}$ ($\vec{c}$) the value profile (cost profile) of all buyers (sellers). Players' utilities are quasi-linear, i.e., buyer i's utility from acquiring a unit at price p_i is $v_i - p_i$, and seller j's utility from selling his unit for payment p_j is $p_j - c_j$. The *welfare* achieved by a set of buyers B and sellers S is the difference between the total value of the $\min\{|B|, |S|\}$-highest buyers and the total cost of the $\min\{|B|, |S|\}$-lowest sellers.

A set of buyers and sellers is *feasible* if the set of buyers is feasible and the set of sellers is feasible and there are at least as many sellers as there are buyers. Which sets of buyers are feasible is expressed as a set system $(N, \mathcal{I}_N)$, where N is the set of all n buyers, and $\mathcal{I}_N \subseteq 2^N$ is a non-empty collection of all the feasible buyer subsets. Similarly, feasible seller sets are given as a set system $(M, \mathcal{I}_M)$, where M is the set of all m sellers, and $\mathcal{I}_M \subseteq 2^M$ is a non-empty collection of all the feasible seller subsets. The set systems that we consider are *accessible*, meaning that for every nonempty feasible set, there exists an element of the set that we can remove such that it remains feasible. An example of an accessible set system is a *downward closed* one, in which for every nonempty feasible set, removing any element of the set results in another feasible set. We assume that the two feasibility set systems are publicly known and can be accessed via feasibility oracles.

2.2. Double Auction Mechanisms

We study direct and deterministic double auction mechanisms, which consist of an *allocation rule* $x(\cdot,\cdot)$ and a *payment rule* $p(\cdot,\cdot)$. The allocation rule takes a pair of value and cost profiles $\vec{v},\vec{c}$ as input, and outputs the set of players who are *accepted*, or *allocated*, for trade. For every buyer i (seller j), $x_i(\vec{v},\vec{c})$ (resp., $x_j(\vec{v},\vec{c})$) indicates whether he is allocated by the mechanism. The payment rule also takes a pair of value and cost profiles $\vec{v},\vec{c}$ as input, and computes payments that it charges the buyers and pays to the sellers. We use $p_i(\vec{v},\vec{c})$ to denote the payment buyer i is charged, and $p_j(\vec{v},\vec{c})$ to denote the payment seller j is paid. A buyer who is not accepted is charged 0 and a seller who is not accepted is paid 0. The *welfare* a mechanism achieves is the welfare of its set of accepted players.

Non-Strategic Properties. We study the following non-strategic properties of double auction mechanisms:

(1) *Feasibility*. A double auction mechanism is feasible if for every value and cost profiles $\vec{v},\vec{c}$, the set of accepted buyers and sellers is feasible. Formally, let B be the set of accepted buyers and let S be the set of accepted sellers, then $B \in \mathcal{I}_N$, $S \in \mathcal{I}_M$ and $|B| \leq |S|$.
(2) *Budget balance (BB)*. A double auction mechanism is budget balanced if for every value and cost profiles $\vec{v},\vec{c}$, the difference between the sum of payments charged from the accepted buyers and the sum of payments paid to the accepted sellers is non-negative.
(3) *Efficiency*. For $\delta \geq 1$, a double auction mechanism is δ-approximately efficient if for every value and cost profiles $\vec{v},\vec{c}$, the welfare it achieves is at least a $(1/\delta)$-fraction of the optimal welfare.

Strategic Properties. We also study the following strategic properties of double auction mechanisms:

(1) *Individual rationality (IR)*. A double auction mechanism is IR if for every value and cost profiles $\vec{v},\vec{c}$, every accepted buyer i is not charged more than his value v_i, and every accepted seller j is paid at least his cost c_j. Non-accepted players are charged/paid zero.
(2) *Dominant strategy incentive compatible (DSIC)*. A double auction mechanism is DSIC if for every value and cost profiles $\vec{v},\vec{c}$ and for every i, j, v'_i, c'_j, it holds that buyer i is (weakly) better off reporting his true value v_i than any other value v'_i, and seller j is (weakly) better off reporting his true cost c_j than any other cost c'_j. Formally,

$$x_i(\vec{v},\vec{c})v_i - p_i(\vec{v},\vec{c}) \geq x_i((v'_i, v_{-i}),\vec{c})v_i - p_i((v'_i, v_{-i}),\vec{c}),$$

and similarly for seller j.
(3) *Weak group-strategyproofness (WGSP)*. A double auction mechanism is WGSP if for every value and cost profiles $\vec{v},\vec{c}$, for every set of buyers and sellers $B \cup S$ and every alternative value and cost reports of these players v'_B, c'_S, there is at least one player in $B \cup S$ who is (weakly) better off when the players report truthfully as when they report v'_B, c'_S. The idea is that such a player does not have a positive incentive to join the deviating group.[1]

[1]A stronger notion of group strategyproofness requires that no group of buyers and sellers can jointly deviate to make some member of the group strictly better off while all other members are no worse off. This stronger notion is violated, for example, if a seller's cost sets the price for a buyer, because then the seller can claim to have a lower cost which will lower the buyer's payment.

The following characterization of DSIC and IR double auction mechanisms follows from standard arguments.

PROPOSITION 2.1. *A double auction mechanism is DSIC and IR if and only if:*

(1) *The allocation rule is* monotone*, i.e., for all value and cost profiles $\vec{v}, \vec{c}$, every accepted buyer who raises his value remains accepted, and every accepted seller who lowers his cost remains accepted.*

(2) *The payment rule applies* threshold payments*, i.e., every accepted buyer is charged his threshold value — the lowest value he could have reported while remaining accepted, and every accepted seller is paid his threshold cost — the highest cost he could have reported while remaining accepted.*

Note that threshold payments are sufficient to guarantee IR, and since all the mechanisms we consider apply threshold payments we can assume individual rationality from now on.

A similarly simple characterization of WGSP and IR double auction mechanisms is not available.[2]

2.3. Running Examples

Matroids. One possible structure of the feasible sets on either side of the market are matroids. A set system $(U, \mathcal{I})$ is a matroid if (1) $\emptyset \in \mathcal{I}$, (2) for all $S \subset T \subseteq U : T \in \mathcal{I}$ implies $S \in \mathcal{I}$ (downward closed property), (3) if $S, T \in \mathcal{I}$ and $|T| > |S|$, then there exists $u \in T \setminus S$ such that $S \cup \{u\} \in \mathcal{I}$ (exchange property). The set U is called the *ground set* and the sets in $\mathcal{I}$ are called *independent*. A maximal independent set is called a *basis*, and a minimal dependent set is called a *circuit*.

Knapsacks. Another possibility is knapsacks. In this case, the elements of the ground set U have publicly-known *sizes* $(s_1, \ldots, s_{|U|})$, and the family of feasible sets $\mathcal{I}$ includes every subset $S \subseteq U$ such that its total size $\sum_{i \in S} s_i$ is at most the capacity C of the knapsack. We denote the ratio between the size of the largest element and the size of the knapsack by $\lambda \leq 1$, and the ratio between the size of the smallest element and the size of the largest element by $\mu \leq 1$.

Matchings. A third class of feasibility restrictions are bipartite matching constraints. In this case the ground set U is the edge set of some bipartite graph $G = (V, U)$, and the family of feasible sets $\mathcal{I}$ are the subsets of the ground set that correspond to bipartite matchings in this graph.

3. COMPOSITION FRAMEWORK

In this section we describe our framework for designing double auctions via composition. The framework includes one-sided algorithms and their composition; accordingly, in order for the resulting double auctions to exhibit good properties, there are two categories of required assumptions, one for the algorithms and the other for the composition. We discuss these in two separate sections: in Section 3.1 we define relevant properties of the algorithms used on each side of the market, and in Section 3.2 we consider different ways of composing these algorithms into double auction mechanisms.

As an example of how the framework works, suppose one cares about designing a double auction with stronger incentives than DSIC. In Section 4 we study WGSP and discover we need the assumptions of deferred-acceptance algorithms on both sides of the market, and of using a composition of backward type and “lookback” composition

[2]For recent progress towards characterizing WGSP and BB mechanisms in the context of cost sharing mechanisms see Juarez [2013].

rule. These assumptions are described in Sections 3.1 and 3.2, respectively. They are also relevant for other double auction properties, in this case BB.

In Section 3.3 we conclude by instantiating our framework for the running examples; these also aim to help in parsing the abstract definitions.

3.1. Ranking Algorithms

The one-sided algorithms we use for our compositions are called *ranking algorithms*. A ranking algorithm for buyers (sellers) is a deterministic algorithm that receives as input a value profile $\vec{v}$ (cost profile $\vec{c}$), and returns an ordered set of buyers (sellers), which we refer to as a *stream*. Not all buyers (sellers) must appear in this stream, e.g., for feasibility considerations (see below). The *rank* of a buyer (seller), denoted by $r_i(\vec{v})$ ($r_j(\vec{c})$), is his position in the stream (e.g., 1 if he appears first), or ∞ if he does not appear in the stream. The streams of players will be traversed greedily, as per our greedy approach. Accessing the next player in the stream is called *querying* the algorithm. The *history* after k queries consists of the identities and values/costs of the k queried players.

Feasibility. We distinguish between two natural feasibility requirements, based on the feasibility set system of the relevant side of the market.

— A *forward-feasible* ranking algorithm returns a stream of players such that any prefix of the stream is a feasible set.
— A *backward-feasible* ranking algorithm returns a stream of players such that there is a minimal prefix of the stream which must be *discarded* to get a feasible set, and such that discarding additional players maintains feasibility. We denote the rank of the first player that does not have to be discarded by ℓ.

The semantic difference between forward-feasible and backward-feasible ranking algorithms is that the former returns a stream of players who can be greedily accepted for trade, while the latter returns a stream of players who can be greedily rejected. This difference can affect the way in which we define the algorithms' other properties, as we now describe.

Consistency. A forward-feasible ranking algorithm is *consistent* if the players appear in decreasing order of attractiveness, i.e., decreasing values for buyers and increasing costs for sellers. A backward-feasible ranking algorithm is consistent if starting from rank ℓ, the players appear in increasing order of attractiveness, i.e., increasing values for buyers and decreasing costs for sellers.

Uniform Approximation. Fix a cardinality t. First consider the welfare-maximizing feasible set of at most t players from each side of the market; denote their total value by $v_{\mathrm{OPT}}(t)$ for buyers, and total cost by $c_{\mathrm{OPT}}(t)$ for sellers. Now for a forward-ranking algorithm, let $v_{\mathrm{ALG}}(t)$ ($c_{\mathrm{ALG}}(t)$) denote the total value (cost) achieved by greedily allocating to the first $\leq t$ buyers (sellers) in the feasible part of the output stream. For a backward-ranking algorithm, the definitions are the same except that the last $\leq t$ buyers or sellers in the feasible part of the output stream are considered.

Let $\alpha, \beta \geq 1$. A ranking algorithm for buyers is a *uniform α-approximation* if for every value profile $\vec{v}$ and every $t \leq n$, it holds that

$$v_{\mathrm{ALG}}(t) \geq \frac{1}{\alpha} \cdot v_{\mathrm{OPT}}(t).$$

A ranking algorithm for sellers is a *uniform β-approximation* if for every cost profile $\vec{c}$ and every $t \leq m$, it holds that

$$c_{\mathrm{ALG}}(t) \leq \beta \cdot c_{\mathrm{OPT}}(t).$$

That is, uniform approximation asserts that every feasible prefix/suffix output by the ranking algorithms is near-optimal under the corresponding cardinality constraint.

Rank Monotonicity. A ranking algorithm is *rank monotone* if for every input and every player, the rank of the player changes monotonically with his bid. More formally, consider a forward-feasible ranking algorithm for buyers. The algorithm is rank monotone if for every value profile $\vec{v}$ and buyer i,

$$v_i' > v_i \implies r_i' \le r_i,$$

where $r_i = r_i(\vec{v})$ is buyer i's rank when he reports value v_i, and $r_i' = r_i((v_i', v_{-i}))$ is his rank when he reports v_i'. Similarly, a forward-feasible ranking algorithm for sellers is rank monotone if for every $\vec{c}$ and seller j,

$$c_j' > c_j \implies r_j' \ge r_j.$$

For a backward-feasible algorithm this implication is reversed.[3]

Deferred-Acceptance. The following definition is obtained by applying the deferred-acceptance definition of Milgrom and Segal [2014] to ranking algorithms. We will use it to generalize the trade reduction mechanism of McAfee, and in the composition theorems.

A *deferred-acceptance ranking algorithm* for buyers (sellers) is a backward-feasible ranking algorithm that computes its output stream as follows. The algorithm maintains a set A of *active* buyers (sellers) — initially all buyers (sellers) — and at each step until this set is empty:

— Scoring functions assign every active buyer (seller) $i \in A$ a non-negative score $s_i^A(v_i, v_{-A})$ (score $s_i^A(c_i, c_{-A})$).[4] The scoring functions are weakly increasing in their first argument.
— The buyer (seller) with the lowest (highest) score is added to the output stream and removed from the set of active players A.

3.2. Composition of Ranking Algorithms

A *composition* between two ranking algorithms is a double auction mechanism, whose allocation rule iteratively queries the ranking algorithms for the next buyer and next seller, and decides whether to accept or reject them based on the current history, and whose payment rule applies threshold payments. The decision whether to accept or reject based on current history is called the *composition rule*.

We focus on two composition types — *forward* and *backward* — which give rise to the *trade exhaustion* and *trade reduction* double auction families.

Forward Composition and Trade Exhaustion. A *forward composition* is a composition between two forward-feasible ranking algorithms, whose allocation rule operates as follows, given a composition rule and value and cost profiles $\vec{v}, \vec{c}$:

(1) Query both ranking algorithms to form a buyer-seller pair (v_i, c_j); if one stream runs out of players go to (3).
(2) Based on the bids/asks seen so far, use the composition rule to decide whether to accept this buyer-seller pair and go to (1) or to go to (3).
(3) Stop, rejecting all remaining players.

[3]I.e., $v_i' > v_i \implies r_i' \ge r_i$ and $c_j' > c_j \implies r_j' \le r_j$.
[4]Note that a player's score is a function of his own bid and the bids of the inactive players, and the scoring function may depend on the set of active players. Importantly, the scoring function may not depend on the bids of other active players.

OBSERVATION 3.1. *Every double auction mechanism obtained from forward composition is feasible.*

A composition rule is *monotone* if for each pair (v_i, c_j) that is accepted given history h, it holds that every other pair $(v_{i'}, c_{j'})$ with value $v_{i'} \geq v_i$ and cost $c_{j'} \leq c_j$ is also accepted given any prefix h' of history h.[5] For example, a "t-threshold" composition rule, which accepts a buyer-seller pair (v_i, c_j) if and only if their gain from trade $v_i - c_j$ exceeds t, is monotone.

A *trade exhaustion* mechanism is a forward composition of consistent, rank monotone ranking algorithms that uses the 0-threshold composition rule.

Backward Composition and Trade Reduction. A *backward composition* is a composition between two backward-feasible ranking algorithms, whose allocation rule operates as follows, given a composition rule and value and cost profiles $\vec{v}, \vec{c}$:

(0) Pre-processing: Query both ranking algorithms and reject players until the remaining sets B of buyers and S of sellers are feasible. Continue to reject players from the long side of the market until $|B| = |S|$.
(1) Query both ranking algorithms to form a buyer-seller pair (v_i, c_j); if the streams run out of players go to (3).
(2) Based on the bids/asks seen so far, use the composition rule to decide whether to reject this buyer-seller pair and go to (1) or to go to (3).
(3) Stop, accepting all remaining players.

OBSERVATION 3.2. *Every double auction mechanism obtained from backward composition is feasible.*

A *lookback* composition rule decides whether to accept or reject a buyer-seller pair *without* taking into account their value and cost; it is only allowed to depend on the values and costs of players previously queried in step (1) of the backward composition. Note that the values and costs of players queried in the preprocessing step (0) are *not* taken into account.[6] For example, a particular lookback rule is the "rejected t-threshold" composition rule, which accepts a buyer-seller pair (v_i, c_j) if and only if there is a buyer-seller pair $(v_{i'}, c_{j'})$ previously rejected in step (1), whose gain from trade $v_{i'} - c_{j'}$ exceeds t.

The *trade reduction* mechanism is a backward composition of consistent, deferred-acceptance ranking algorithms that uses the rejected 0-threshold rule.

Remark 3.3. While trade reduction mechanisms are inherently based on backward composition, it is possible to implement trade exhaustion mechanisms as forward or backward compositions, in the sense that whether the implementation is forward or backward is irrelevant for our composition theorems. The implication of a forward versus backward implementation is that it dictates which ranking algorithms must be used — forward-feasible or backward-feasible. In the applications we consider in this paper, forward-feasible ranking is arguably more natural, and so we focus on trade exhaustion based on forward composition.

3.3. Running Examples

We give a high-level description of ranking algorithms for our running examples, focusing on the maximization (buyer-side) versions of the problems. A detailed description appears in the full version of the paper.

---

[5]For a backward composition, this needs to hold for every history h' such that h is a prefix of h'.

[6]Note that a lookback composition rule is trivially monotone as the decision in step (2) to start accepting pairs only depends on previously seen pairs.

Matroids. The greedy algorithm for matroids sorts the elements of the ground set by their weight, from highest to lowest, and then goes through the list of elements adding the next element if its addition does not violate feasibility. It gives rise to a consistent, rank monotone, and a uniform 1-approximation [Edmonds 1971].

It is not difficult to see that the order in which the elements are processed can be turned around. For this go through the elements in reverse order, and reject an element if it forms a circuit with a subset of the unrejected elements. In fact, this property is a structural property of the active bidders (i.e., no knowledge of their bids is required). The reverse greedy algorithm can therefore be implemented as a deferred-acceptance algorithm.

Knapsacks. The standard greedy algorithm for knapsacks ranks elements by density, which leads to a non-consistent rule. An alternative greedy algorithm, which is often also a good heuristic, sorts the elements by their weight — from high to low — and keeps adding elements to the knapsack as long as they fit into the knapsack. This is clearly consistent and rank monotone. The standard argument for bounding the performance of this algorithm relates its performance to the degree to which the knapsack is filled and the degree by which it can be off from the weight density. A straightforward extension of this argument shows that for all cardinalities k its outcome is within a factor of $(1-\lambda)\mu$ of the optimal solution of cardinality k.

Again it is possible to reverse this greedy algorithm. For this we simply have to go through the list of elements from low to high, and reject the next element if the elements that are still to come do not fit in the knapsack. This backward greedy algorithm has the same properties as the forward greedy algorithm. Since the "remaining elements fit into the knapsack" property is again a structural property of the active bidders alone (i.e., it does not depend on their bids), it can be implemented as a deferred-acceptance algorithm.

Matchings. For bipartite matchings the greedy algorithm that sorts edges by their weight, and accepts the next edge if none of its endpoints was previously added gives rise to a consistent, rank monotone, uniform 2-approximation [Hassin and Rubinstein 2002]. This can be improved upon by computing a maximum weight matching of the edge weights squared. The resulting algorithm is a consistent, rank monotone, uniform $\sqrt{2}$-approximation [Hassin and Rubinstein 2002].

Unfortunately, it is not clear how to reverse these algorithms — the obvious attempt to traverse the elements in reverse order and to reject the next edge only if none of its endpoints is covered by edges still to come fails miserably. It turns out that a different approach that first grows a path of locally heaviest edges and then selects a subset of these edges is a consistent, rank monotone, uniform 2-approximation that can be implemented as a deferred-acceptance algorithm.

4. INCENTIVES

In this section we present a DSIC and a WGSP composition theorem. We then apply these theorems to trade exhaustion and trade reduction.

4.1. DSIC Composition Theorem

Our composition theorem for DSIC applies equally well to forward and backward composition. One might suspect that monotonicity of all involved components — the composition rule and the ranking algorithms — is sufficient for DSIC, but an example deferred to the full version of the paper shows that consistency is needed as well. It is also shown in the full version that under a weak "no waste" assumption, rank monotonicity is necessary for achieving a DSIC double auction.

THEOREM 4.1. *A double auction that is obtained from forward composition of consistent, rank monotone ranking algorithms for the buyers and the sellers using a monotone composition rule is DSIC.*

PROOF. We apply the characterization of DSIC double auctions in Proposition 2.1 to show that the above composition is DSIC. The composition applies threshold payments by definition, and so all we need to show is that its allocation rule is monotone. Fix value and cost profiles $\vec{v}, \vec{c}$. We show monotonicity by arguing that an accepted buyer who raises his value remains accepted (a similar argument shows that an accepted seller who lowers his cost remains accepted).

Denote the accepted buyer by i. Denote the seller with whom i trades by j. Let h be the history when the composition rule is applied to (i, j). By rank monotonicity of the buyer ranking algorithm, if i raises his value then his rank weakly decreases. Let j' be the seller with which i is considered for trade after his rank decreases. Then $c'_j \leq c_j$ by consistency of the seller ranking algorithm. Since the composition rule is monotone, and the history h' when it is applied to (i, j') is a prefix of h, then the pair (i, j') must be accepted for trade as well. □

4.2. WGSP Composition Theorem

Our WGSP composition theorem shows that the lookback composition of deferred-acceptance ranking algorithms is WGSP. To prove this we show that the composition can itself be implemented as a deferred-acceptance algorithm. The theorem then follows from the incentive properties of deferred-acceptance algorithms that Milgrom and Segal [2014] establish.

THEOREM 4.2. *A double auction that is obtained from backward composition of deferred-acceptance ranking algorithms for the buyers and the sellers using a lookback composition rule is WGSP.*

The proof of this theorem, which is given in the full version of the paper, works by compiling the scoring functions for the buyers and sellers, available from the deferred-acceptance ranking algorithms, into a common scoring for all players; a deferred-acceptance implementation of the composition then follows directly. The challenge is in producing a valid scoring that implements the composition rule — as Section 6 will imply, this is not possible for all composition rules, and the proof must therefore utilize the properties of the lookback composition.

Moreover, we shall see that the proof must fully utilize the deferred-acceptance framework of Milgrom and Segal [2014]. The common scoring will not depend only on the player's own report, but critically also on the reports of already rejected players. This is in contrast to all deferred-acceptance algorithms we are aware of that precede [Milgrom and Segal 2014] — consider for example the work of Bikhchandani et al. [2011], as well as many other works cited therein.

4.3. Incentives of Trade Exhaustion and Trade Reduction

Two immediate corollaries of our composition theorems are incentive guarantees for trade exhaustion and trade reduction.

COROLLARY 4.3. *The trade exhaustion double auction is DSIC.*

COROLLARY 4.4. *The trade reduction double auction is WGSP.*

Remark 4.5. Two further implications of our reduction are that the trade reduction mechanism can be implemented as a clock auction (from Proposition 13 in [Milgrom and Segal 2014]), and that there exists a Nash equilibrium of the trade reduction mechanism with first price payments, in which the allocation and payments are identical to the DSIC outcome of the trade reduction mechanism with threshold payments (from Proposition 20).

Remark 4.6. It should be noted that none of the stronger incentive properties of the trade reduction mechanism are shared by the trade exhaustion mechanism: It is not implementable as a clock auction, it is not weakly group-strategyproof, and there does not always exist an outcome-equivalent Nash equilibrium in the corresponding first price auction.

Remark 4.7. Milgrom and Segal also show additional incentive properties for a restricted class of deferred-acceptance algorithms that are "non-bossy," meaning that players can only alter the set of accepted players if they change their own status from accepted to rejected or vice versa. Trade reduction — through its dependence on rejected bids — violates non-bossiness: the buyer-seller pair that is prevented from trade can misreport, claiming to have negative gain from trade, thus leading to the elimination of another buyer-seller pair that would otherwise have been accepted.

4.4. Running Examples

Our DSIC composition theorem and the corollary for trade exhaustion apply to all forward algorithms discussed in Section 3.3. The WGSP composition theorem and the corollary for trade reduction apply to all backward algorithms discussed in Section 3.3.

5. WELFARE

In this section we discuss the welfare guarantees of double auction mechanisms arising from compositions. We present a welfare composition theorem, and use it to derive welfare guarantees for trade exhaustion and trade reduction.

5.1. Welfare Composition Theorem

The welfare composition theorem proves a bound that depends on the number of trades t in the double auction, the number of trades t' that would have maximized the gain from trade in the double auction, and the number of trades t^* in the optimal solution. The number of trades t and t' may differ e.g. as a result of trade reduction; t' and t^* may differ e.g. when the ranking algorithms are approximate.

As is standard with mixed-sign objective functions (c.f., [Roughgarden and Sundararajan 2009]), the bound also depends on the optimal gain from trade via $\gamma \geq 1$ which we define such that

$$v_{\mathrm{OPT}}(t^*) = \gamma \cdot c_{\mathrm{OPT}}(t^*).$$

In the welfare-maximizing solution, the total value is certainly at least the cost; $\gamma \geq 1$ is the factor by which it is bigger. Intuitively, the closer γ is to 1, the closer the optimal welfare is to 0, and the harder it is to get a relative approximation of the optimal welfare.

THEOREM 5.1. *Consider a double auction that is obtained by composing consistent ranking algorithms for the buyers and the sellers that are uniform α- and β-approximations. Assume further that the composition rule only accepts buyer-seller pairs (i, j) with $v_i - c_j \geq 0$. Let t, t', and t^* as well as γ be defined as above. Then for $\gamma > \alpha\beta \geq 1$ we have*

$$v_{\mathrm{ALG}}(t) - c_{\mathrm{ALG}}(t) \geq \frac{t}{t'} \cdot \frac{\frac{\gamma}{\alpha} - \beta}{\gamma - 1} \cdot \big(v_{\mathrm{OPT}}(t^*) - c_{\mathrm{OPT}}(t^*)\big).$$

Note that if $\alpha = \beta = 1$, then the second term on the right hand side of the inequality vanishes. In this case the approximation factor is simply the fraction of optimal trades made. The bound degrades gracefully from this ideal case in all relevant parameters.

PROOF. Since the double auction is composed of consistent ranking algorithms we can number the buyers and sellers in the respective streams so that $v_1 \geq v_2 \geq \cdots \geq v_n$ and $c_1 \leq c_2 \leq \cdots \leq c_m$. Using this notation,

$$v_{\text{ALG}}(t) - c_{\text{ALG}}(t) = \sum_{i=1}^{t}(v_i - c_i) \text{ and } v_{\text{ALG}}(t') - c_{\text{ALG}}(t') = \sum_{i=1}^{t'}(v_i - c_i).$$

Another implication of the fact that the double auction is composed of consistent ranking algorithms is that the gain from trade is non-increasing. That is, $i < j$ implies $v_i - c_i \geq v_j - c_j$. Hence for all s such that $t < s \leq t'$ we have $v_s - c_s \leq \frac{1}{t}\sum_{i=1}^{t}(v_i - c_i)$. It follows that

$$\begin{aligned} v_{\text{ALG}}(t) - c_{\text{ALG}}(t) &= \sum_{i=1}^{t'}(v_i - c_i) - \sum_{i=t+1}^{t'}(v_i - c_i) \\ &\geq \sum_{i=1}^{t'}(v_i - c_i) - (t' - t)\frac{1}{t}\sum_{i=1}^{t}(v_i - c_i) \\ &= \big(v_{\text{ALG}}(t') - c_{\text{ALG}}(t')\big) - \left(\frac{t'}{t} - 1\right)\big(v_{\text{ALG}}(t) - c_{\text{ALG}}(t)\big). \end{aligned}$$

Rearranging this shows

$$v_{\text{ALG}}(t) - c_{\text{ALG}}(t) \geq \frac{t}{t'}\big(v_{\text{ALG}}(t') - c_{\text{ALG}}(t')\big). \quad (1)$$

Next we use the fact that by the definition of t' all trades up to and including t' are beneficial, while all subsequent trades yield a deficit. That is, $v_s - c_s \geq 0$ for $s \leq t'$ and $v_s - c_s < 0$ for $s > t'$. Hence,

$$v_{\text{ALG}}(t') - c_{\text{ALG}}(t') \geq v_{\text{ALG}}(t^*) - c_{\text{ALG}}(t^*). \quad (2)$$

Finally, we use that the ranking algorithms are uniform α- and β-approximations and the definition of γ to deduce that

$$\begin{aligned} v_{\text{ALG}}(t^*) - c_{\text{ALG}}(t^*) &\geq \frac{1}{\alpha}v_{\text{OPT}}(t^*) - \beta c_{\text{OPT}}(t^*) \\ &= \left(\frac{\gamma}{\alpha} - \beta\right)c_{\text{OPT}}(t^*) \\ &= \frac{\frac{\gamma}{\alpha} - \beta}{\gamma - 1}\big(v_{\text{OPT}}(t^*) - c_{\text{OPT}}(t^*)\big). \end{aligned} \quad (3)$$

To prove the claim it now suffices to combine inequality (1), inequality (2), and inequality (3). □

5.2. Welfare of Trade Exhaustion and Trade Reduction

Next we use the general welfare composition theorem to derive welfare guarantees for the trade exhaustion and trade reduction rules. The claimed bounds follow directly from the observation that trade exhaustion makes all t' possible trades, while trade reduction makes exactly $t' - 1$ trades. Note that it suffices to show the bound for $\gamma > 1$, as for $\gamma = 1$ the optimal welfare is zero.

COROLLARY 5.2. *Suppose $\gamma > 1$ and that the ranking algorithms used in the trade exhaustion double auction are uniform α- and β-approximations. Then the approximation ratio is at least*

$$\frac{\frac{\gamma}{\alpha} - \beta}{\gamma - 1}.$$

COROLLARY 5.3. *Suppose $\gamma > 1$ and that the ranking algorithms used in the trade reduction double auction are uniform α- and β-approximations. Then the approximation ratio is at least*

$$\left(1 - \frac{1}{t'}\right)\left(\frac{\frac{\gamma}{\alpha} - \beta}{\gamma - 1}\right).$$

5.3. Running Examples

The greedy-by-weight algorithms for matroids and knapsacks as well as the three algorithms for matchings discussed in Section 3.3 satisfy the conditions of the welfare composition theorem and its corollaries. For matroids this means that there is a welfare-maximizing DSIC trade exhaustion mechanism, and a $(1 - 1/t^*)$-approximate WGSP trade reduction mechanism.

6. INCENTIVES AND WELFARE

In this section we investigate the interplay between incentives and welfare. We prove a lower bound that applies to WGSP mechanisms, and an impossibility result for mechanisms based on forward composition. The lower bound for WGSP mechanisms shows that for settings in which consistent, uniformly optimal deferred-acceptance algorithms are available, the trade reduction double auction based on these algorithms is optimal subject to WGSP.

6.1. Lower Bound Subject to WGSP

Our lower bound for WGSP mechanisms applies to anonymous mechanisms; mechanisms whose outcome does not change if the names of the players are permuted. Recall that t^* denotes the number of trades in the optimal solution.

THEOREM 6.1. *No anonymous double auction mechanism that is WGSP can achieve an approximation ratio strictly better than $1 - \frac{1}{t^*}$.*

A simple observation is that the welfare-maximizing mechanism is not WGSP: Consider a setting with one buyer with value v, and one seller with cost c. Since the buyer has to win if and only if $v \geq c$, it has to pay c. A similar argument for the seller shows that it must be paid v. Now the buyer and seller can jointly deviate from truthful reports $v \geq c$ to falsified reports $v' > v \geq c > c'$, which leads to a strict decrease/increase of their payments to v' and c'.

The proof of Theorem 6.1, which is given in the full version of the paper, is complicated by the fact that it has to deal with approximation mechanisms and there is typically more than one way to achieve a certain approximation guarantee. In fact, it is not difficult to see that any approximation guarantee between 1 and $1 - 1/t^*$ can be supported by a DSIC mechanism. Our lower bound shows that only the extreme point $1 - 1/t^*$ is also achievable by a WGSP mechanism.

6.2. Impossibility Result for Forward Composition

We now show an impossibility result — no double auction mechanism that is obtained from forward composition can be WGSP and achieve an approximation ratio of $1 - 1/t^*$. In fact in the following proposition we prove a stronger impossibility result, which shows that forward composition is particularly ill-equipped to achieve either WGSP or budget balance while maintaining a non-trivial efficiency guarantee. The proof is given in the full version of the paper.

PROPOSITION 6.2. *For every DSIC double auction obtained from forward composition of consistent ranking algorithms, there are value and cost profiles such that one of the following holds:*

(*1*) *The budget deficit is arbitrarily high and the mechanism is not WGSP.*
(*2*) *The welfare is arbitrarily smaller than the welfare achieved by the trade reduction double auction.*

7. BUDGET BALANCE

In this section we first establish that the trade reduction mechanism is BB. Previously, this was only known for the simple setting with no feasibility constraints other than having an equal number of buyers and sellers. Afterwards we establish two lower bounds on the welfare achievable by a DSIC and BB mechanism, one for the unconstrained setting which applies to all mechanism and on for the constrained setting which applies to double auction mechanisms resulting from composition.

7.1. Budget Balance Composition Theorem

We begin by showing that the trade reduction double auction is BB. The proof exploits a property of deferred-acceptance algorithms, namely that a winning player can only alter the set of winning players by becoming a loser.

THEOREM 7.1. *The trade reduction double auction is BB.*

PROOF. The trade reduction double auction is composed of two backward-feasible ranking algorithms. Rename the players such that the values ordered by rank are $(v_1, \ldots, v_n)$ and the costs ordered by rank are $(c_1, \ldots, c_n)$. Let ℓ_B (ℓ_S) be the rank of the first buyer (seller) who does not have to be rejected to obtain feasibility. Let ℓ' be the rank of the buyer-seller pair reduced by trade reduction. It holds that $\ell' \geq \max\{\ell_B, \ell_S\}$.

Our goal will be to show that the buyers accepted by the trade reduction double auction pay at least the value $v_{\ell'}$. A symmetric argument shows that the accepted sellers are paid at most the cost $c_{\ell'}$. This is sufficient to establish budget balance since $v_{\ell'} \geq c_{\ell'}$. Since the double auction uses threshold payments, it is enough to show that any accepted buyer $i > \ell'$ who lowers his value to $v'_i < v_{\ell'}$ will no longer be accepted.

Let r' denote i's rank after lowering his value to v'_i. Assume first that $r' < \ell_B$. Since the scores by which the deferred-acceptance algorithm ranks buyers depend only on the buyers' own values and on values of previously rejected buyers, the set of buyers ranked 1 to $r' - 1$ does not change following the decrease in i's value. By definition of ℓ_B, rejecting buyers 1 to $r' - 1$ is not enough to make the remaining buyers a feasible set, and so buyer i who is ranked r' is necessarily rejected.

Now assume $r' \geq \ell_B$. By consistency, buyers with rank at least ℓ_B are ranked in increasing order of their values. This order is the same as before i's value decreased, except that i now appears before buyer ℓ'. Therefore, the buyer-seller pair i now belongs to either has negative gain from trade or is reduced; in both cases buyer i is rejected. □

7.2. Lower Bounds Subject to DSIC and BB

Our first theorem regarding optimality subject to DSIC and BB, whose proof is given in the full version of the paper, gives a lower bound on the approximation ratio of DSIC and BB double auctions; this lower bound is detail-free and complements lower bounds on the welfare achievable by BIC and BB mechanisms in Bayesian settings [Satterthwaite and Williams 2002].

THEOREM 7.2. *No DSIC and BB double auction can achieve an approximation ratio strictly better than* $1 - \frac{1}{t^*}$.

Our second theorem, whose proof is again given in the full version of the paper, gives a lower bound on the welfare achievable by a double auction resulting from composition imposing DSIC and BB.

THEOREM 7.3. *Consider a double auction that is obtained by composing consistent, rank monotone ranking algorithms for the buyers and the sellers, which are uniform* α*- and* β*-approximations. Then subject to DSIC and BB, the approximation ratio of this double auction cannot be better than*

$$\left(1 - \frac{1}{t'}\right)\left(\frac{\frac{\gamma}{\alpha} - \beta}{\gamma - 1}\right).$$

8. CONCLUSION

In this paper we proposed a modular approach to the design of double auctions that decomposes the design task into the task of designing greedy algorithms for either side of the market and a composition rule. Focusing on the unit-demand and unit-supply case, we proved a number of composition theorems for (approximate) efficiency, DSIC or WGSP, and BB that relate the properties of the double auction to the properties of the modules used in its construction.

A challenging direction for future work is to extend our approach to even more complex settings, such as:

— Settings with identical goods but multi-unit demand or multi-unit supply, say with downward-sloping valuations or upward-sloping costs.
— Settings with non-identical goods in which each seller has a single good to sell and either single-minded buyers or unit-demand buyers.

ACKNOWLEDGMENTS

We greatly benefitted from discussion with Moshe Babaioff, Gabriel Caroll, Fuhito Kojima, Paul Milgrom and Ilya Segal. This research is supported in part by NSF Awards CCF-1016885 and CCF-1215965, an ONR PECASE Award, an SNF Postdoctoral Fellowship, and a Hsieh Family Stanford Interdisciplinary Graduate Fellowship.

REFERENCES

M. Babaioff and N. Nisan. Concurrent auctions across the supply chain. *Journal of Artifical Intelligence Research*, 21:595–629, 2004.

M. Babaioff, N. Nisan, and E. Pavlov. Mechanisms for a spatially distributed market. *Games and Economic Behavior*, 66(2):660–684, 2009.

C. Y. Baldwin and K. B. Clark. *Design Rules: The Power of Modularity*. MIT Press, 2000.

S. Bikhchandani, S. de Vries, S. J. Schummer, and R. Vohra. An ascending Vickrey auction for selling bases of a matroid. *Operations Research*, 59(2):400–413, 2011.

L. Blumrosen and S. Dobzinski. Reallocation mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation*, 2014. To appear.

J. Bredin and D. C. Parkes. Models for truthful online double auctions. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence*, pages 50–59, 2005.

G. Carroll. The efficiency-incentive tradeoff in double auction environments. Working paper, 2013.

L. Y. Chu. Truthful bundle/multiunit double auction. *Management Science*, 55(7): 1184–1198, 2009.
M. W. Cripps and J. M. Swinkels. Efficiency of large double auctions. *Econometrica*, 74(1):47–92, 2006.
P. Dütting, V. Gkatzelis, and T. Roughgarden. The performance of deferred-acceptance auctions. In *Proceedings of the 15th ACM Conference on Economics and Computation*, 2014. To appear.
J. Edmonds. Matroids and the greedy algorithm. *Mathematical Programming*, 1:127–136, 1971.
D. Fudenberg, M. Mobius, and A. Szeidl. Existence of equilibrium in large double auctions. *Journal of Economic Theory*, 133(1):550–567, 2007.
M. Gonen, R. Gonen, and E. Pavlov. Generalized trade reduction mechanisms. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 20–29, 2007.
R. Hassin and S. Rubinstein. Robust matchings. *SIAM Journal on Discrete Mathematics*, 15:530–537, 2002.
L. Hurwicz. On informationally decentralized systems. In C. B. McGuire and R. Radner, editors, *Decision and Organization*, pages 297–336. University of Minnesota Press, Minneapolis, MN, USA, 1972.
R. Juarez. Group strategyproof cost sharing: The role of indifferences. *Games and Economic Behavior*, 82:218–239, 2013.
F. Kojima and T. Yamashita. Double auction with interdependent values: Incentives and efficiency. Mimeo, 2013.
R. P. McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2): 434–450, 1992.
P. Milgrom and I. Segal. Deferred-acceptance auctions and radio spectrum reallocation. In *Proceedings of the 15th ACM Conference on Economics and Computation*, 2014. To appear.
A. Mu'alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64:612–631, 2008.
R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.
T. Roughgarden and M. Sundararajan. Quantifying inefficiency in cost-sharing mechanisms. *Journal of the ACM*, 56(4):article 23, 2009.
A. Rustichini, M. A. Satterthwaite, and S. R. Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5):1041–63, 1994.
M. A. Satterthwaite and S. R. Williams. The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. *Review of Economic Studies*, 56(4):477–98, 1989.
M. A. Satterthwaite and S. R. Williams. The optimality of a simple market mechanism. *Econometrica*, 70(5):1841–1863, 2002.