Efficient Two-Sided Markets with Limited Information*

Paul Dütting
Google Research
Zürich, Switzerland
duetting@google.com

Federico Fusco
Department of Computer, Control,
and Management Engineering,
Sapienza University of Rome
Rome, Italy
fuscof@diag.uniroma1.it

Philip Lazos
Department of Computer, Control,
and Management Engineering,
Sapienza University of Rome
Rome, Italy
lazos@diag.uniroma1.it

Stefano Leonardi
Department of Computer, Control,
and Management Engineering,
Sapienza University of Rome
Rome, Italy
leonardi@diag.uniroma1.it

Rebecca Reiffenhäuser
Department of Computer, Control,
and Management Engineering,
Sapienza University of Rome
Rome, Italy
rebeccar@diag.uniroma1.it

ABSTRACT

A celebrated impossibility result by Myerson and Satterthwaite (1983) shows that any truthful mechanism for two-sided markets that maximizes social welfare must run a deficit, resulting in a necessity to relax welfare efficiency and the use of approximation mechanisms. Such mechanisms in general make extensive use of the Bayesian priors. In this work, we investigate a question of increasing theoretical and practical importance: how much prior information is required to design mechanisms with near-optimal approximations?

Our first contribution is a more general impossibility result stating that no meaningful approximation is possible without any prior information, expanding the famous impossibility result of Myerson and Satterthwaite.

Our second contribution is that one *single sample* (one number per item), arguably a minimum-possible amount of prior information, from each seller distribution is sufficient for a large class of two-sided markets. We prove matching upper and lower bounds on the best approximation that can be obtained with one single sample for subadditive buyers and additive sellers, regardless of computational considerations.

Our third contribution is the design of computationally efficient blackbox reductions that turn any one-sided mechanism into a two-sided mechanism with a small loss in the approximation, while using only one single sample from each seller. On the way, our

blackbox-type mechanisms deliver several interesting positive results in their own right, often beating even the state of the art that uses full prior information.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic game theory and mechanism design**; • **Applied computing** → **Economics**.

KEYWORDS

mechanism design with samples, two-sided markets, truthful mechanisms, social welfare, budget balance

ACM Reference Format:

Paul Dütting, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. 2021. Efficient Two-Sided Markets with Limited Information. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC '21)*, June 21–25, 2021, Virtual, Italy. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3406325.3451076>

1 INTRODUCTION

Two-sided markets play an increasingly important role in practice, with applications ranging from internet trading platforms (such as eBay), to ridesharing platforms (such as Uber or Lyft), to ad exchanges (such as AppNexus). A canonical two-sided market consists of sellers that provide certain services or goods to the market, and buyers that consume these services or goods. It is natural to assume that both sides of the market have private information about their production costs and willingness to pay.

The design of such two-sided markets with private information has a long tradition in economics [33, 43]. One typically strives to find mechanisms that guarantee individual rationality (IR), dominant strategy incentive compatibility (DSIC), budget balance (BB) (i.e., that the mechanism collects more money from the buyers than it gives to the sellers), all while also maximizing social welfare. Designing such mechanisms, however, turns out to be rather challenging. Consider, for example, the celebrated Vickrey-Clarke-Groves (VCG) mechanism [12, 25, 43]. This mechanism chooses an allocation that maximizes social welfare, and charges each agent the amount by which the agent's presence reduces the social welfare of the other agents. Already Vickrey [43] observed that when applied

*Part of the work of Paul Dütting was done while he was an Associate Professor in the Department of Mathematics at London School of Economics. The work of Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser was supported in part by ERC Advanced Grant 788893 AMDROMA "Algorithmic and Mechanism Design Research in Online Markets" and MIUR PRIN project ALGADIMAR "Algorithms, Games, and Digital Markets."

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '21, June 21–25, 2021, Virtual, Italy

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8053-9/21/06.

<https://doi.org/10.1145/3406325.3451076>

to two-sided markets, this mechanism is IR, DSIC, and maximizes social welfare, but may produce a deficit.

EXAMPLE 1 (TWO-SIDED VCG FOR A SINGLE GOOD). *Consider a seller owning a single indivisible good, which she values at v_s , and two interested buyers willing to pay v_{b_1} and v_{b_2} , respectively. Assume $v_s < v_{b_1} > v_{b_2}$, so that the socially efficient outcome is to give the item to b_1 , achieving a welfare of v_{b_1} . The VCG mechanism will assign the item to b_1 , whose participation keeps the item from being assigned to s or b_2 , so she will pay $\max\{v_s, v_{b_2}\}$. For the seller s , in her absence no one can get the item and welfare is reduced by v_{b_1} , so the mechanism will pay her this amount. Recall that $\max\{v_s, v_{b_2}\} - v_{b_1} < 0$, so the mechanism runs a deficit.*

This example is actually fairly robust, extending to settings with Bayesian priors on multiple buyers and multiple sellers. A famous impossibility result of Myerson and Satterthwaite [33] in fact shows that this is not just a problem of VCG, it is a more fundamental incompatibility: any mechanism that is IR and DSIC and maximizes social welfare must violate budget balance even if the mechanism knows the Bayesian priors of the buyers and sellers. In this paper, we explore the fundamental information-theoretic as well as computational requirements for achieving near-optimal IR, DSIC, and BB mechanisms.

1.1 Information-Theoretic Barrier: Necessity of Some Prior Dependence

Since information about the underlying distributions of the agents' valuations is often scarce, it would be ideal to find an approximation mechanism that guarantees close to optimal welfare without the use of such prior information. However, this is impossible: extending the famous inefficiency theorem of Myerson and Satterthwaite [33], we prove that without prior information, the welfare cannot be approximated up to any constant:

THEOREM 1 (IMPOSSIBILITY; SECTION 3). *No mechanism M for bilateral trade (or richer variants of it) without prior information, where agents' valuations for a single item take arbitrary values in $[0, \infty)$, can be IC (in expectation), IR, BB and also guarantee an α -approximation to the optimal social welfare, for any fixed value $\alpha > 1$.*

As mentioned above, this crucial necessity of prior information is bad news: even when there is some information available on the priors, it usually presents itself as a limited amount of incomplete data and assuming to indeed know the underlying distributions exactly is often unrealistic.

1.2 Information-Theoretically Optimal Mechanisms

In light of our new impossibility result: (Theorem 1), an important question is how much information about the priors is needed to enable approximately welfare-maximizing IR, DSIC, and BB mechanisms. Various forms of limited access to the priors are conceivable. A standard vehicle in Algorithmic Game Theory—and the lens we adopt here—is sample complexity (e.g., [3, 13, 18, 19, 31, 38]). That is, we assume that the mechanism designer has access to samples, possibly from both the buyers' and sellers' distributions, and can use this knowledge when setting up the mechanism.

Our first result is that for the very general class of *subadditive* buyers and additive sellers, a *single sample* from each seller (and *no samples* from the buyers) suffices for a factor 2 approximation to the optimal social welfare.

To this end we consider the following *adjusted-objective VCG mechanism*: it takes a single sample from each seller distribution, and offers the sellers the sample as price. Afterwards, it runs the (one-sided) VCG on the buyers and items of sellers who accepted their price. But instead of considering the usual objective of welfare maximization, it selects an allocation and payments based on the adjusted welfare (= total value of the buyers minus sampled value of the respective items).¹

THEOREM 2 (ADJUSTED VCG; SECTION 4). *For subadditive buyer valuations and unit-supply/additive sellers, the social welfare output by the adjusted-objective VCG mechanism which uses a single sample from each seller and no samples from the buyers yields a 2-approximation to the optimal welfare, in expectation. The mechanism is IR, DSIC, and BB.*

This result is in fact fairly robust, and generally needs only a certain form of “approximate location oracle” for the seller distributions: a single sample is a particular economic way to obtain such an oracle, but other aggregate statistics such as the median or other quantiles (possibly at the expense of worse approximation guarantees) also suffice.

Theorem 2 yields poly-time mechanisms for settings where the adjusted VCG can be computed efficiently. This includes gross substitutes valuations for the buyers [35], so we obtain an IR, DSIC, BB 2-approximation for this class of valuations with a single sample from each seller. Another special case is bilateral trade, where our theorem implies the existence of a 2-approximate IR, DSIC, and BB single-sample mechanism.

We complement our positive result concerning the adjusted VCG mechanism with the following lower bound, which shows that the factor 2 is best possible for mechanisms that—just as the adjusted VCG mechanism—are *deterministic*, in that the only randomness in these mechanisms stems from the randomness in the sample they receive.

THEOREM 3 (LOWER BOUND; SECTION 4). *Every deterministic IC, IR and BB mechanism for bilateral trade that receives as sole prior information a single sample from the seller's distribution has approximation ratio at least 2.*

We note that [8, 9] previously showed a factor 2-approximate IR, DSIC, and BB mechanism for bilateral trade—the *median mechanism*—and that no deterministic mechanism that uses only information from one side of the market can yield a better approximation.² For more general valuations, the state of the art is a 6-approximation for XOS valuations [14, 15] and a 8-approximation for subadditive valuations [8].

¹ This idea of a penalty modifying the objective of the VCG mechanism is also used by Blumrosen and Dobzinski [8] (however with the median instead of a sample), which directly applied to our problems results in an 8-approximation given subadditive buyer valuations, and with some modification to the proof one can show a factor of 4.

² Our lower bound is similar, but it does not immediately follow from their result. The reason is that mechanisms that use only one sample cannot be simulated by a deterministic mechanism using prior information.

We extend the factor 2 of [9] for bilateral trade to single-sample mechanisms, and the much more general problem with subadditive buyers and additive sellers.

While these results nicely illuminate the information-theoretical situation, the approach of adjusting the objective of a one-sided mechanism to accommodate seller payments has important limitations. Essentially, it can be applied only via (generalized) VCG because of its properties as an *affine maximizer*, but VCG is not necessarily computationally feasible and cannot be used, e.g., in online environments.

1.3 Poly-Time Near-Optimal Single-Sample Mechanisms

We next address the shortcomings and limited flexibility of the adjusted objective approach. In general, it is not even clear what it means “to run a given one-sided approximation mechanism on the adjusted objective.” We will instead use that for several subclasses of subadditive (such as XOS or GS) we can adjust buyer valuations by deducting item prices, and the resulting valuations will still be within that same class. We leverage this as follows: As in our approach in Section 1.2 we use a single sample from each seller distribution. We offer this sampled price to the sellers. Then we run the given one-sided approximation mechanism on the buyers and the items of those sellers who accepted their sampled price, with buyer valuations reduced by the respective item prices. We charge the buyers whatever the one-sided mechanism would charge them plus the sum of the sampled prices of the items they receive; the sellers whose items were assigned to buyers receive the sampled price.

The constructed mechanism is BB by design. We show that if the one-sided mechanism is DSIC then so is the two-sided mechanism. From the approximation guarantee of the one-sided mechanism we get a guarantee with respect to the buyer surplus (w.r.t. the original valuations and sampled prices), and we show that this translates (approximately) to the actual objective we are interested in. Moreover, if the one-sided mechanism is online, then so is the two-sided mechanism (for the buyers).

THEOREM 4 (BLACK BOX I; SECTION 5). *Denote by α the approximation guarantee of any one-sided IR, DSIC offline/online mechanism for maximizing social welfare for XOS valuations. We give a two-sided mechanism for XOS buyers and unit-supply sellers that is IR, DSIC, BB, uses a single sample from each seller and provides a $\max\{2\alpha, 3\}$ approximation to the optimal social welfare. The two-sided mechanism inherits the offline/online properties of the one-sided mechanism on the buyer side and is offline on the seller side.*

Two implications of Theorem 4 are: for matching constraints (unit demand valuations), we can apply the poly-time truthful secretary algorithm of [37]; this yields a IR, DSIC, and BB $2e$ -approximate single-sample mechanism that is online random order on the buyer side and offline on the seller side. For general XOS functions we obtain a poly-time IR, DSIC, and BB $O((\log \log m)^3)$ -approximate single-sample mechanism by plugging in the state of the art truthful approximation mechanism for this class [1, 2, 20]. Both of these are the first single-sample mechanisms for the respective settings. They compare to a 6-approximation which can process buyers in any order, but requires additional information [14, 15].

Our second reduction is more restricted since it applies only to unit-demand buyers and unit-supply sellers with identical items, but it improves on the first reduction in two ways: it yields a strongly BB mechanism in which the total payments collected from the buyers exactly match those to the sellers, and it additionally allows online random arrival of the sellers.

The reduction runs the given one-sided mechanism to select buyers and tentative payments for the buyers, it draws a single sample from each seller distribution, and then randomly matches the selected buyers and sellers offering them to trade at the maximum of the two prices.

THEOREM 5 (BLACK BOX II; SECTION 6). *Denote by α the approximation guarantee of any one-sided IR, DSIC offline/online mechanisms for the intersection of a downward closed set system $\mathcal{I}$ with a uniform matroid that may or may not use information on the priors. We give a two-sided mechanism for unit-demand buyers, unit-supply sellers with identical goods, and constraints $\mathcal{I}$ on the buyers, that is IR, DSIC, SBB, has the same information requirements on the buyer side as the one-sided mechanism and uses a single sample from each seller, and yields a $(1 + 1/(2 - \sqrt{3}) \cdot \alpha) \approx (1 + 3.73 \cdot \alpha)$ approximation to the optimal social welfare. The mechanism inherits the online/offline properties of the one-sided mechanism on the buyer side and it is online random order on the seller side.*

Theorem 5 leads to additional results for double auctions. We specifically obtain a $1 + 3.73(1 + o(1))$ -approximate mechanism for a k -uniform matroid constraint on the buyers by using the truthful secretary algorithm of [27]. The mechanism admits random order arrivals of buyers and sellers, and requires a single sample on the seller side. We also obtain a $1 + 3.73(1 + o(1))$ -approximate mechanism for a k -uniform matroid with fixed order arrivals of the buyers and random order arrivals from the sellers by using the single sample prophet inequality algorithm of [3]. This second mechanism requires a single sample from each buyer and seller. (All our other mechanisms only use samples from the sellers.)

Note that besides this variety of results for different valuation classes, arrival models, and even prior information assumptions on the buyer side, any future contribution regarding the one-sided problems will also yield a new result for the two-sided version via our constructions.

1.4 Related Work

A landmark result in economics/mechanism design is Myerson and Satterthwaite’s impossibility result [33], which states that there can be no IC (BIC or DSIC) mechanism that is budget balanced and maximizes social welfare. The same paper also describes the “second best” mechanism which maximizes social welfare subject to the other constraints.

In a sequence of papers Satterthwaite et al. [39–41] considered the non-truthful *buyer’s bid mechanism* for double auction settings with identical items, and showed that with i.i.d. buyers and i.i.d. sellers as the number of buyers and sellers $n, m \rightarrow \infty$ equilibrium bids converge to truthful bids and the corresponding allocation converges to an efficient allocation. In a very elegant paper, McAfee [29] proposed the DSIC and budget balanced *trade reduction mechanism* for double auction settings with identical items. This mechanism allows all but the least efficient trade. Thus, when ℓ is the number

of trades in an efficient solution, it achieves a $1 - 1/\ell$ approximation. McAfee's trade reduction idea has since been extended to a variety of considerably more general settings (see, e.g., [5, 6, 21]).

A pioneering paper by Blumrosen and Dobzinski [8] proposed the *median mechanism* for bilateral trade, which sets a posted price equal to the median of the seller's distribution and shows that this mechanism obtains a 2-approximation to the optimal welfare. Subsequent work by the same authors [9] has shown that this factor 2 is actually optimal for deterministic mechanisms that only use information about the seller distribution, and improved the approximation guarantee to $e/(e-1)$ through a randomized mechanism whose prices depend on the seller distribution in a more intricate way. Colini-Baldeschi et al. [14] gave a fixed-price mechanism that achieves a $25/13 \approx 1.92$ approximation, and show a lower bound of ≈ 1.3360 that applies to any DSIC and strongly budget balanced mechanism. Most recently, Kang and Vondrák [26] gave a $e/(e-1) - \epsilon$ approximate mechanism for bilateral trade.

Colini-Baldeschi et al. [14, 15] also develop DSIC and strongly budget balanced mechanisms for more general settings with combinatorial buyers and unit-supply sellers, using ideas from prophet inequalities and posted-price mechanisms for one-sided markets. The first paper gives a 16-approximate DSIC and strongly budget balanced mechanism for double auction settings where the buyers can be subject to a matroid constraint. This mechanism is online fixed order on the buyer side, and online random order on the seller side. The second paper gives a 6-approximate DSIC and strongly budget balanced mechanism for XOS buyers and unit-supply sellers that is online fixed order on the buyer side and offline on the seller side.

A parallel line of work [4, 10, 11, 16, 30, 42] has considered the related but different objective of optimizing the *gain from trade*, which measures the expected increase in total value that is achievable by applying the mechanism, with respect to the initial allocation to the sellers. Gain from trade is harder to approximate than social welfare, and $O(1)$ approximations of the optimal Bayesian mechanism are only possible in BIC implementations.

Goldner et al. [23] recently suggested an alternative, resource augmentation approach to gains from trades in two-sided markets. They ask how many buyers (resp. sellers) need to be added into the market so that a variant of McAfee's trade reduction mechanism yields a gain from trade superior to the optimal gains from trade in the original market. As a side product they obtain a 4-approximate single-sample mechanism for gains from trade under natural conditions on the distributions.

Another related line of work considers the problem of maximizing revenue in double auction settings, either in static environments [24, 34] or in dynamic environments [7]. A variation where both buyers and sellers arrive dynamically and the mechanism can hold on to items was investigated in [22, 28]. On a higher level, our work is related to prior work that has shown how having a single sample enables the design of mechanisms that achieve near-optimal revenue (e.g., [19]) or prophet inequalities and posted-price mechanisms with near optimal welfare (e.g., [3, 17, 18]). It is also related to prior work on composing mechanisms in an incentive compatible manner, such as [32] for one-sided markets and [21] for two-sided markets.

2 MODEL AND DEFINITIONS

Two-Sided Markets. In a two-sided market we are given a set of n buyers $\mathcal{B}$ and a set of m sellers $\mathcal{S}$. Each seller s is *unit supply*, i.e., she has a single indivisible item for sale, and a private valuation $v_s \in \mathbb{R}_{\geq 0}$ for the item she sells. Each buyer b has a private valuation function $v_b : 2^{\mathcal{S}} \rightarrow \mathbb{R}_{\geq 0}$, mapping each set of sellers to a non-negative real. All results hold also for the equivalent case of sellers with additive valuations. We write $v_{\mathcal{B}}$ and $v_{\mathcal{S}}$ for the vector of buyer valuations and seller valuations, respectively. Buyer and seller valuations are drawn independently from distributions F_b for $b \in \mathcal{B}$ and F_s for $s \in \mathcal{S}$.

The valuation functions of the buyers will be constrained to come from some class of functions $\mathcal{V}$. Buyers are *unit demand* if for each buyer b and set of sellers S , $v_b(S) = \max_{s \in S} \{v_b(\{s\})\}$. Buyers have *subadditive* valuations if for each b and sets of sellers S_1, S_2 , $v_b(S_1 \cup S_2) \leq v_b(S_1) + v_b(S_2)$. Buyers have *fractionally subadditive* (or XOS) valuations if for each buyer b and every set of sellers S , $v_b(S) = \max_{a \in A_b} \sum_{s \in S} a(s)$, where A_b is a set of additive valuation functions. We say that items are *identical* if the valuation function $v_b(S)$ of all buyers b only depends on the cardinality of the set S they receive, i.e., for all b and all $S, S' \subseteq \mathcal{S}$ with $|S| = |S'|$ we have $v_b(S) = v_b(S')$. Otherwise, items are *non-identical*.

We also allow for constraints on which buyers can trade simultaneously. We express these constraints through set systems $\mathcal{I}_{\mathcal{B}} \subseteq 2^{\mathcal{B}}$. We require these set systems $\mathcal{I}_{\mathcal{B}}$ to be downward closed. That is, whenever $X \subseteq Y$ and $Y \in \mathcal{I}_{\mathcal{B}}$, then also $X \in \mathcal{I}_{\mathcal{B}}$. Of particular interest for our work will be *matroids*, i.e., downward-closed set systems that additionally satisfy a natural exchange property. Formally: whenever $A, B \in \mathcal{I}$ and $|A| > |B|$ then there exists $a \in A \setminus B$ such that $B \cup \{a\} \in \mathcal{I}$. A special case are *k-uniform matroids* where $A \in \mathcal{I}$ whenever $|A| \leq k$.

An *allocation* is a partition of the sellers $\mathcal{S}$ into n disjoint sets $(S_1, \dots, S_n)$, i.e., $\bigcup_i S_i \subseteq \mathcal{S}$ and $S_i \cap S_j = \emptyset$ for all $i \neq j$, with the interpretation that buyer b_i for $1 \leq i \leq n$ receives the items of the sellers in S_i . An allocation $A = (S_1, \dots, S_n)$ is *feasible* if the set of buyers $\mathcal{B}_A = \{b_i \in \mathcal{B} \mid S_i \neq \emptyset\}$ that receive a non-empty allocation is admissible (i.e., $\mathcal{B}_A \in \mathcal{I}_{\mathcal{B}}$). The *social welfare* of an allocation $A = (S_1, \dots, S_n)$ is given by the sum of the valuations that buyers b_i for $1 \leq i \leq n$ have for the items of the sellers in their respective sets S_i plus the valuations of the sellers that are not assigned to any buyer, i.e.,

$$\text{SW}(A) = \text{SW}(S_1, \dots, S_n) = \sum_{b_i \in \mathcal{B}} v_{b_i}(S_i) + \sum_{s \in \mathcal{S}, s \notin \bigcup_i S_i} v_s.$$

We use $\text{OPT}(v_{\mathcal{B}}, v_{\mathcal{S}})$ to denote the feasible allocation that maximizes social welfare.

Mechanisms. A (direct revelation) *mechanism* $M = (x, p)$ receives bids $\text{bid}_b : 2^{\mathcal{S}} \rightarrow \mathbb{R}_{\geq 0}$ from each buyer $b \in \mathcal{B}$ and $\text{bid}_s \in \mathbb{R}_{\geq 0}$ from each seller $s \in \mathcal{S}$. The bids of the buyers are constrained to be consistent with the class of functions $\mathcal{V}$ of their valuations. Bids represent reported valuations, and need not be truthful. In analogy to our notation for valuations, we use $\text{bid}_{\mathcal{B}}$ and $\text{bid}_{\mathcal{S}}$ for the vector of bids of all buyers or all sellers, respectively.

A mechanism M is defined through an *allocation rule* $x : \mathcal{V}^n \times \mathbb{R}_{\geq 0}^m \rightarrow \times_{i=1}^n 2^{\mathcal{S}}$ and a *payment rule* $p : \mathcal{V}^n \times \mathbb{R}_{\geq 0}^m \rightarrow \mathbb{R}^{n+m}$.

The mapping from bids to feasible allocations can be randomized, in which case $x(\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}})$ is a random variable. The payments can also be randomized. We interpret the vector of payments as the payments that the buyers need to make to the mechanism, and that the sellers receive from the mechanism. We use the shorthand $p_b(\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}})$ and $p_s(\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}})$ to refer to the payment from buyer b to the mechanism and from the mechanism to seller. We require that $p_s = 0$ if seller s keeps her item.

A mechanism is *single-sample* if the only information it is given about the two sets of distributions F_b for $b \in \mathcal{B}$ and F_s for $s \in \mathcal{S}$ is a single sample from each of these distributions.

Utilities. We assume that buyers and sellers have *quasi-linear utilities*, and that they are utility maximizers. It means that the utility $u_b^M(v_b, (\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}}))$ of buyer b with valuation function v_b in mechanism $M = (x, p)$ under bids $(\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}})$ is given by her valuation for the items she receives minus payment. That is,

$$u_b^M(v_b, (\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}})) = \mathbb{E}[v_b(x_i(\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}})) - p_b(\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}})],$$

where the expectation is over the randomness in the mechanism. The utility $u_s^M(v_s, (\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}}))$ of a seller s in mechanism $M = (x, p)$ under bids $(\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}})$ is the payment she receives if she sells the item and her valuation for her item otherwise. Formally, $u_s^M(v_s, (\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}}))$ is equal to

$$\mathbb{E}[\mathbb{I}_{\{s's \text{ item is sold}\}} \cdot p_s(\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}}) + \mathbb{I}_{\{s's \text{ item is not sold}\}} \cdot v_s],$$

where the expectation is over the randomness in the mechanism.

Goals. We seek to design mechanisms, and specifically single-sample mechanisms, with the following desirable properties:

(1) **Individual Rationality.** Mechanism $M = (x, p)$ is *individually rational (IR)* if for all $b \in \mathcal{B}$ $u_b^M(v_b, (v_{\mathcal{B}}, v_{\mathcal{S}})) \geq 0$ and for all $s \in \mathcal{S}$, $u_s^M(v_s, (v_{\mathcal{B}}, v_{\mathcal{S}})) \geq v_s$.

(2) **Incentive Compatibility.** Mechanism $M = (x, p)$ is (*dominant-strategy*) *incentive compatible (DSIC)* or *truthful* if for each buyer b and each seller s , all valuation functions v_b and v_s , all possible bids $\text{bid}_{\mathcal{B}}$ by the buyers, and all possible bids by the sellers $\text{bid}_{\mathcal{S}}$, it holds that

$$\begin{aligned} u_b^M(v_b, ((v_b, \text{bid}_{\mathcal{B} \setminus \{b\}}, \text{bid}_{\mathcal{S}})) &\geq u_b^M(v_b, (\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}})), \\ u_s^M(v_s, (\text{bid}_{\mathcal{B}}, (v_s, \text{bid}_{\mathcal{S} \setminus \{s\}}))) &\geq u_s^M(v_s, (\text{bid}_{\mathcal{B}}, \text{bid}_{\mathcal{S}})), \end{aligned}$$

where $\text{bid}_{\mathcal{B} \setminus \{b\}}$ denotes the set of all the buyer bids but b 's and $\text{bid}_{\mathcal{S} \setminus \{s\}}$ denotes the set of all the seller bids but s 's.

(3) **Budget Balance.** A truthful mechanism $M = (x, p)$ is *weakly budget balanced (BB or WBB)* if

$$\mathbb{E} \left[\sum_{b \in \mathcal{B}} p_b(v_b, v_{\mathcal{S}}) - \sum_{s \in \mathcal{S}} p_s(v_{\mathcal{B}}, v_s) \right] \geq 0,$$

and it is *strongly budget balanced (SBB)* if the above holds with equality.

(4) **Efficiency.** Finally, a truthful mechanism $M = (x, p)$ provides an α -approximation to the optimal social welfare, for some $\alpha \geq 1$, if it holds that $\alpha \cdot \mathbb{E}[\text{SW}(x(v_{\mathcal{B}}, v_{\mathcal{S}}))] \geq \mathbb{E}[\text{SW}(\text{OPT}(v_{\mathcal{B}}, v_{\mathcal{S}}))]$.

Remark. Our mechanisms will actually satisfy even stronger IR and BB properties in that they will satisfy these conditions “ex post” (i.e., pointwise).

3 IMPOSSIBILITY OF APPROXIMATION WITHOUT PRIOR INFORMATION

In this section we show that at least a tiny amount of information about the priors is needed in order to obtain meaningful approximations of the social welfare with IR, DSIC, and BB mechanisms.

We present here the proof of [Theorem 1](#) for deterministic mechanisms, which nicely illustrates the difficulties at the core of this statement. The considerably more technical, detailed proof for the general theorem can be found in [Appendix A](#).

THEOREM 1 (IMPOSSIBILITY). *No mechanism M for bilateral trade (or richer variants of it) without prior information, where agents' valuations for a single item take arbitrary values in $[0, \infty)$, can be IC (in expectation), IR, WBB and also guarantee an α -approximation to the optimal social welfare, for any fixed value $\alpha > 1$.*

PROOF. Fix any constant approximation ratio α , and assume that M is an IR, weakly budget-balanced, truthful α -approximation for the problem. For the ease of exposition assume M is deterministic: the same proof strategy works in the most general case.

Assume now that $v_b > \alpha \cdot v_s$. The only way M can guarantee this ratio is by trading the item from s to b . Doing this truthfully and in an individually rational fashion without running a deficit implies that for the price p_b charged from the buyer, and p_s paid to the seller, it holds that $v_s \leq p_s \leq p_b \leq v_b$.

Recall that intuitively, the only play here is to set $p_s = v_b$ and $p_b = v_s$. In fact, for truthfulness, an agent should not be able to influence the price she pays, and when determining said price, the only value we therefore know for sure will be in the feasible interval is the reported valuation of the *opposite* agent. Our argument is close to this intuition, however a bit more intricate.

We assumed M to be truthful, i.e., if v_s is a fixed value, then b 's utility will be at least as good when she reports the true v_b as when reporting some v'_b . Since otherwise, whenever v_b is actually the value with the higher price, the agent is incentivized to make a false report. Note that M is forced to make a trade for both of the values in order to preserve an α -approximation. We have established so far that the buyer price p_b^* will be a single value for all buyer reports in $(\alpha \cdot v_s, \infty)$, and for all of those buyer reports, there also needs to be a trade. By budget balance, when this is the case it holds $p_b^* \in [v_s, \alpha \cdot v_s]$ which means that even the smallest buyer for which M has to trade will be willing to accept p_b^* .

This shows that whenever the seller reports a fixed value v_s , the price for each *critical* buyer will come from a fixed interval starting at v_s .

We can do the same from the other direction. Start fixing some buyer report v_b ; since M is an α -approximation, it will have to make a trade whenever $v_s \in [0, v_b/\alpha]$. Because M is truthful also for the sellers, each seller report in this interval must result in the same price p_s^* . To be acceptable for all sellers in the critical interval, similarly as above, it must be that $p_s^* \in [v_b/\alpha, v_b]$. We put this together with the fact that, due to weak budget balance, p_s^* cannot exceed p_b^* , i.e., $[v_b/\alpha, v_b] \ni p_s^* \leq p_b^* \in [v_s, \alpha \cdot v_s]$. This goes awry

for any choice of numbers p_s^*, p_b^* once v_s and v_b are such that the possible intervals for p_s^*, p_b^* do not overlap. Setting for example $v_s^* = 1$ and $v_b^* = \alpha^3$ violates the previous inequality. $\square$

4 TWO-APPROXIMATION FOR SUBADDITIVE BUYERS

In this section, we present a single-sample, IR, BB and DSIC 2-approximation for subadditive buyers. Our mechanism is essentially the generalized VCG mechanism, which we use as a blackbox.

Given a sample v'_s for each seller, we do the following:

- For each $s \in S$, offer a price v'_s to s and if she accepts, add s to the set $\hat{S}$ of available items.
- On the instance $(B, \hat{S})$, run VCG with the objective of maximizing the affine function $\sum_{b \in B} v_b(S_b) - \sum_{s \in \bigcup_{b \in B} S_b} v'_s$, where S_b denotes the items assigned to b .
- Assign the items in $\bigcup_{b \in B} S_b$ to their respective buyers, then charge them the price determined by VCG and pay each according seller a price of v'_s . All other sellers keep their item.

Note that VCG outputs an optimal assignment according to above objective because of its well-known properties as an affine maximizer. Showing that our mechanism is DSIC, IR, and BB is not too hard and has been done before, see [8] for more details. Intuitively it is IR because we ask the sellers if they want to trade and for the buyers, VCG is of course IR. It is DSIC for the sellers since they can refuse, and for the buyers because VCG is DSIC.

Finally, it is BB because for every buyer b , the sum of her VCG payments equals at least the sum of her items' samples, which is what we pay to the sellers. This is true because the loss introduced to the other agents by b 's presence in the VCG routine is at least the sum of samples of items S_b : by not assigning those at all in a solution without b , we gain exactly this amount according to our adjusted objective.

Proving that the mechanism is also a 2-approximation is considerably more involved and will need some preparation. For simplicity, we use notation for an allocation and its welfare interchangeably. Let OPT_{max} denote an allocation on (B, S) which maximizes the adjusted objective

$$\sum_{b \in B} \left(v_b(S_b^{OPT_{max}}) - \sum_{s \in S_b^{OPT_{max}}} \max\{v_s, v'_s\} \right).$$

Further, for any allocation A over sets (S, B) with valuations v_s, v_b , respectively, its corresponding welfare is $\sum_{s \in S} v_s + \sum_{b \in B} T_b^A$, where $T_b^A = v_b(A(b)) - \sum_{s \in A(b)} v_s$. For use in the proof, we further define the *adjusted* buyer gains $\hat{T}_b^A = v_b(A(b)) - \sum_{s \in A(b)} \max\{v_s, v'_s\}$.

In order to prove the approximation, we first show three preliminary statements which we then combine. The first one, [Lemma 1](#), reports a folklore insight on subadditive functions and sampling. A proof is added for completeness.

LEMMA 1. *Let $f : 2^N \rightarrow \mathbb{R}$ be a subadditive set function, and let X be a random subset of the base set N where each element appears independently with probability $\frac{1}{2}$, then $\mathbb{E}[f(X)] \geq \frac{1}{2}f(N)$.*

PROOF. Every subset X' of N is realized with the same identical probability $p = 2^{-|N|}$. Therefore we have the following:

$$\begin{aligned} \mathbb{E}[f(X)] &= \sum_{X' \subseteq N} p \cdot f(X') = \frac{1}{2} \sum_{X' \subseteq N} p \cdot (f(X') + f(N \setminus X')) \\ &\geq \frac{1}{2} \sum_{X' \subseteq N} p \cdot f(N) = \frac{1}{2} f(N). \end{aligned}$$

Notice that the inequality is given by subadditivity. $\square$

The other two lemmas we need are more technical and specific to our problem. [Lemma 2](#) states that the loss in buyer value when replacing seller valuations with $\max\{v_s, v'_s\}$ is not too large, i.e., the buyers' contribution to OPT_{max} is closely related to their contribution to OPT .

LEMMA 2. *The following inequality holds true:*

$$\sum_{b \in B} \mathbb{E}[\hat{T}_b^{OPT_{max}}] \geq \sum_{b \in B} \mathbb{E}[T_b^{OPT}] - \sum_{s \in S} \mathbb{E}[(v'_s - v_s)_+]$$

PROOF. Define the set of items assigned to any buyer by the optimal assignment OPT to be $S^{OPT} \subseteq S$. Let $B^{max} \subseteq B$ be those buyers for which their endowments $S_b^{OPT} \subseteq S^{OPT}$ satisfy

$$v_b(S_b^{OPT}) > \sum_{s \in S_b^{OPT}} \max\{v_s, v'_s\}$$

and call the set of items sold by OPT to these buyers S^{max} . Consider the assignment A^+ implied by sets B^{max}, S^{max} . This assignment clearly is feasible under the adjusted objective, which we can model for the purpose of finding an assignment by replacing all seller valuations with $\max\{v_s, v'_s\}$ in VCG. Recall we denote with $\hat{T}_b$ the according, adjusted buyer gains T_b . We have

$$\sum_{b \in B \setminus B^{max}} \hat{T}_b^{A^+} = 0, \text{ while } \sum_{b \in B \setminus B^{max}} T_b^{OPT} \leq \sum_{s \in S \setminus S^{max}} (v'_s - v_s)_+.$$

For the other items, i.e., those in S^{max} , we have

$$\begin{aligned} \sum_{b \in B^{max}} \hat{T}_b^{A^+} &= \sum_{b \in B^{max}} v_b(S_b^{OPT}) - \sum_{s \in S^{max}} \max\{v_s, v'_s\} \\ &= \sum_{b \in B^{max}} T_b^{OPT} - \sum_{s \in S^{max}} (v'_s - v_s)_+. \end{aligned}$$

Putting both together, this yields

$$\sum_{b \in B} \hat{T}_b^{A^+} \geq \sum_{b \in B} T_b^{OPT} - \sum_{s \in S} (v'_s - v_s)_+$$

and since OPT_{max} maximizes the sum of $\hat{T}_b$, also $\sum_{b \in B} \hat{T}_b^{OPT_{max}} \geq \sum_{b \in B} \hat{T}_b^{A^+}$. $\square$

Our next lemma, [Lemma 3](#), will allow us to relate our adjusted VCG mechanism to OPT_{max} . Let us denote the output of adjusted-objective VCG as ALG .

LEMMA 3. *The following inequality holds true:*

$$\mathbb{E} \left[\sum_{b \in B} T_b^{ALG} \right] \geq \frac{1}{2} \mathbb{E} \left[\sum_{b \in B} \hat{T}_b^{OPT_{max}} \right]$$

PROOF OF LEMMA 3. Fix buyer b and her assigned set $S_b^{OPT_{max}}$. Let for each b the assignment $S_b^{1/2}$ be constructed by erasing each $s \in S_b^{OPT_{max}}$ with probability $1/2$, independently from everything else. It holds by definition

$$\hat{T}_b^{OPT_{max}} = v_b(S_b^{OPT_{max}}) - \sum_{s \in S_b^{OPT_{max}}} \max\{v_s, v'_s\}.$$

With subadditivity and Lemma 1, this implies for the assignment defined by the $S_b^{1/2}$ and the according $\hat{T}_b^{1/2}$:

$$\mathbb{E}[\hat{T}_b^{1/2}] \geq \frac{1}{2}v_b(S_b^{OPT_{max}}) - \frac{1}{2} \sum_{s \in S_b^{OPT_{max}}} \max\{v_s, v'_s\}$$

since every item is simply erased with fixed probability $1/2$, in both the left and the right hand side. Note that the expectation is only with respect to the erasing process, while the valuations are fixed.

Let ALG' be the solution OPT_{max} restricted to those items for which $v_s \leq v'_s$, which is the case for each s with probability at least $1/2$, independently. Note that since OPT_{max} only cares about $\max\{v_s, v'_s\}$, it is oblivious to which value realizes the maximum. We infer $\mathbb{E}[\hat{T}_b^{ALG'}] \geq \mathbb{E}[\hat{T}_b^{1/2}]$ for each b , because ALG' can be seen as simply fixing a certain way of flipping the coins for $S_b^{1/2}$, and then retaining some extra items (alternatively, we could employ a random tie breaking). Next, we note that every subset $S_b^{ALG'}$ assigned to a buyer in ALG' is contained in $\hat{S}$, and for the objective of our algorithm, assigning it is advantageous, i.e., yields nonnegative contribution. ALG might deviate from these assignments, but it holds that

$$\sum_{b \in B} v_b(S_b^{ALG}) - \sum_{s \in \bigcup_{b \in B} S_b^{ALG}} v'_s \geq \sum_{b \in B} v_b(S_b^{ALG'}) - \sum_{s \in \bigcup_{b \in B} S_b^{ALG'}} v'_s$$

simply because ALG maximizes the objective. Also, due to our assumption $v_s \leq v'_s$ for all items assigned by ALG' , we get that the right side is equal to $\sum_{b \in B} \hat{T}_b^{ALG'}$, yielding

$$\sum_{b \in B} v_b(S_b^{ALG}) - \sum_{s \in \bigcup_{b \in B} S_b^{ALG}} v'_s \geq \sum_{b \in B} \hat{T}_b^{ALG'}.$$

Further, since ALG only assigns items in $\hat{S}$, i.e., with $v_s \leq v'_s$, we can replace the left side in the same way and get

$$\sum_{b \in B} \hat{T}_b^{ALG} \geq \sum_{b \in B} \hat{T}_b^{ALG'}.$$

Now taking the expectations, it holds

$$\mathbb{E}\left[\sum_{b \in B} \hat{T}_b^{ALG}\right] \geq \mathbb{E}\left[\sum_{b \in B} \hat{T}_b^{ALG'}\right] \geq \mathbb{E}\left[\sum_{b \in B} \hat{T}_b^{1/2}\right]$$

because for each b , above we showed $\mathbb{E}[\hat{T}_b^{ALG'}] \geq \mathbb{E}[\hat{T}_b^{1/2}]$. Plugging in our bound $\mathbb{E}\left[\sum_{b \in B} \hat{T}_b^{1/2}\right] \geq \frac{1}{2}\mathbb{E}\left[\sum_{b \in B} \hat{T}_b^{OPT_{max}}\right]$ implies

$$\mathbb{E}\left[\sum_{b \in B} \hat{T}_b^{ALG}\right] \geq \mathbb{E}\left[\frac{1}{2} \sum_{b \in B} \hat{T}_b^{OPT_{max}}\right]$$

Finally, we conclude the proof by observing that

$$\sum_{b \in B} T_b^{ALG} \geq \sum_{b \in B} \hat{T}_b^{ALG}. \quad \square$$

With these lemmas at hand, we can now prove the claimed approximation guarantee.

THEOREM 2 (ADJUSTED VCG). *For subadditive buyer valuations, the social welfare output by the adjusted-objective VCG mechanism yields a 2-approximation to the optimal welfare, in expectation. Moreover, the mechanism is IR, DSIC, and BB.*

PROOF. Note that we are interested in the welfare of the allocation, not its adjusted objective. It holds

$$\begin{aligned} \mathbb{E}[ALG] &= \sum_{b \in B} \mathbb{E}[T_b^{ALG}] + \sum_{s \in S} \mathbb{E}[v_s] \\ &\geq \frac{1}{2} \mathbb{E}\left[\sum_{b \in B} \hat{T}_b^{OPT_{max}}\right] + \sum_{s \in S} \mathbb{E}[v_s] \\ &\geq \frac{1}{2} \sum_{b \in B} \mathbb{E}[T_b^{OPT}] - \frac{1}{2} \sum_{s \in S} \mathbb{E}[(v'_s - v_s)_+] + \sum_{s \in S} \mathbb{E}[v_s] \\ &\geq \frac{1}{2} \sum_{b \in B} \mathbb{E}[T_b^{OPT}] + \frac{1}{2} \sum_{s \in S} \mathbb{E}[v_s] = \frac{1}{2} \mathbb{E}[OPT]. \end{aligned}$$

Here, for the first inequality we used Lemma 3, and for the second Lemma 2. Finally, the last inequality is true because for each seller s , it holds that $\mathbb{E}[(v'_s - v_s)_+] \leq \mathbb{E}[v_s]$. In fact, $(v'_s - v_s)_+ \leq v'_s$ and v_s and v'_s share the same distribution. $\square$

We conclude this section by showing that the result in Theorem 2 is best-possible in the sense that no deterministic mechanism with one seller sample as sole prior information can achieve a better approximation.

THEOREM 3. *Every deterministic IC, IR, and BB mechanism for bilateral trade that receives as sole prior information a single sample from the seller's distribution has approximation ratio at least 2.*

PROOF. Because the mechanism is DSIC, the reported bid of the seller cannot affect *how much* she is paid for a given outcome. Her price can however depend on the sample and the buyer's report. By DSIC, WBB and IR we also know that the price the buyer needs to pay whenever she gets the item does not depend on her report and needs to be higher than the seller price but below her own valuation. Using this we construct different cases where the optimum is to always trade, but show that the mechanism cannot do so in all of them. First, we assume that there exist a sample $\hat{s}$ and a buyer report $b \geq 2\hat{s} + \epsilon$ such that the price paid to the seller for a trade is greater than $b \geq 2\hat{s} + \epsilon$. This suggests that the mechanism overcharges: indeed, by taking advantage of DSIC, WBB and IR we can show that a buyer with value $2\hat{s} + \epsilon/2$ would miss the trade. This implies that the mechanism doesn't charge more than twice the seller sample when a trade *has* to be made to guarantee the 2-approximation. We can therefore fix the buyer value to be very high and construct a probability distribution for the seller ranging over many small, exponentially decreasing values. Since the mechanism cannot increase the price too much (given the sample observed), with probability approaching 0.5 the seller will reject the trade.

Let's now formalize the above argument. Let v_b and v_s denote the random variables corresponding to the seller and the buyer valuations and let $p_s(s, b)$ be the price paid to the seller in the event

of a trade when the mechanism receives sample $s \sim F_S$ and report b from the buyer. We distinguish two cases:

Case 1: $\exists \hat{s}, b, \epsilon > 0$ such that $p_s(\hat{s}, b) \geq 2 \cdot \hat{s} + \epsilon$ and $b \geq 2 \cdot \hat{s} + \epsilon$. In this case consider a seller whose value is always exactly $\hat{s}$ and a buyer whose value is exactly $2 \cdot \hat{s} + \epsilon/2$. Clearly, a trade should always happen and the optimal welfare is $2 \cdot \hat{s} + \epsilon/2$. Assuming that trade happens: by weak budget balance and individual rationality, the price $p_b(\hat{s}, 2 \cdot \hat{s} + \epsilon/2)$ paid by the buyer must satisfy:

$$\hat{s} \leq p_s(\hat{s}, 2 \cdot \hat{s} + \epsilon/2) \leq p_b(\hat{s}, 2 \cdot \hat{s} + \epsilon/2) \leq 2 \cdot \hat{s} + \epsilon/2.$$

However, by DSIC and given that a trade also needs to happen for $v_b = b$, we have that:

$$p_b(\hat{s}, b) \geq 2 \cdot \hat{s} + \epsilon > 2 \cdot \hat{s} + \epsilon/2 \geq p_b(\hat{s}, 2 \cdot \hat{s} + \epsilon/2).$$

This violates the DSIC condition, as the buyer can obtain a better price by misreporting, leading to a contradiction. Therefore, one of the two trades doesn't occur and the approximation is at least 2.

Case 2: $\forall \hat{s}, b \quad p_s(\hat{s}, b) \leq 2 \cdot \hat{s}$ or $b \leq 2 \cdot \hat{s}$. Consider a seller whose valuation v_s takes a uniformly random value in $\{1/3, 1/3^2, \dots, 1/3^k\}$ and a buyer with fixed value $v_b = 1$. Clearly, the optimal welfare is 1. Notice that the probability of a trade happening is:

$$\begin{aligned} \mathbb{P}(p_s(\hat{s}, 1) \geq v_s) &= \sum_{i=1}^k \mathbb{P}\left(p\left(\frac{1}{3^i}, 1\right) \geq v_s\right) \cdot \mathbb{P}\left(\hat{s} = \frac{1}{3^i}\right) \\ &= \sum_{i=1}^k \mathbb{P}\left(p\left(\frac{1}{3^i}, 1\right) \geq v_s\right) \cdot \frac{1}{k} \\ &\leq \sum_{i=1}^k \frac{k-i}{k} \cdot \frac{1}{k} = \frac{k-1}{2 \cdot k} \end{aligned}$$

where we used that $p(1/3^i, 1) < 1/2^{i-1}$ therefore the price posted does not reach the next possible valuation. Clearly, this converges to a 2-approximation as k grows to infinity. $\square$

5 COMPUTATIONALLY EFFICIENT COMBINATORIAL MECHANISMS

In this section we study the SURPLUS MECHANISM, which achieves the property described in Theorem 4. In Section 5.1 we present the algorithm and in Section 5.2 we show how it achieves the claimed properties. Finally, in Section 5.3 we discuss computational aspects.

5.1 The Surplus Mechanism

The basic idea behind our SURPLUS MECHANISM (Algorithm 1) is to run the given truthful one-sided mechanism M_α on *discounted buyer valuations* and on *a subset of the sellers*. Note that the problem can be viewed as finding a hypermatching in a bipartite hypergraph $G = (\mathcal{B} \cup \mathcal{S}, E, v_{\mathcal{B}})$ with hyperedge set E defined as all tuples (b, S) s.t. $b \in \mathcal{B}, S \subseteq \mathcal{S}$.

First, given valuations v_S and samples v'_S for each seller, we determine a subset of the sellers $\hat{S}$ as follows. For each $s \in \mathcal{S}$ we put s in $\hat{S}$ if $v_s \leq v'_s$. Otherwise, we will drop s from our considerations. Next we determine discounted valuations. For a given buyer b and a given set of sellers $S \subseteq \mathcal{S}$ let $a_{b,S}$ denote the additive supporting function of buyer b for set S . We define the discounted valuation

that buyer b has for the set of sellers $S \subseteq \hat{S}$ as:

$$\hat{v}_b(S) = \sum_{s \in \hat{S}} (a_{b,S}(s) - v'_s), \text{ where } \hat{S} = \operatorname{argmax}_{S^* \subseteq \mathcal{S} \cap \hat{S}} \left\{ v_b(S^*) - \sum_{s \in S^*} v'_s \right\}. \quad (1)$$

Adjusting the valuations like this retains the XOS property of the original valuations, as shown in the following Proposition.

PROPOSITION 1. *The adjusted valuations defined in Equation (1) are XOS.*

PROOF. The definition of $\hat{v}_b$ in Equation (1) captures the fact that only items in $\hat{S}$ are available for trade, and whenever a buyer is assigned some item s , she will be required to pay an additional v'_s later. Therefore, her valuation for S is no more than her valuation for the best subset of it in $\hat{S}$, minus the sum of the according v'_s . This function is XOS because it can be described as the maximum over the following XOS-support: for all a from the support of the original v_b and $s \in \mathcal{S}$, define $\hat{a}(s) = (a(s) - v'_s)_+$ if $s \in \hat{S}$, and 0 otherwise.

For the $\bar{v}$, defined as

$$\bar{v}_b(S) = \left(\sum_{s \in \mathcal{S} \cap \hat{S}} (a_{b,S}(s) - v'_s) \right)_+,$$

we define the XOS support $\bar{a}(s) = (a(s) - v'_s)$ if $s \in \hat{S}$, and 0 otherwise, and add an additional function a_0 which is 0 for any $s \in \mathcal{S}$. $\square$

Analogously, one can prove that the same holds for the gross substitutes class: after the adjustment, GS valuations remain GS.

Given these *closure* properties, we run the one-sided mechanism M_α on the resulting hypergraph $\hat{G} = (\mathcal{B} \cup \hat{\mathcal{S}}, \hat{E}, \hat{v}_{\mathcal{B}})$ consisting of all buyers, only the sellers in $\hat{S}$, and hyperedge valuations $\hat{v}_{\mathcal{B}}$. This will lead to an allocation $S_1, \dots, S_n$ and payments $p_b^{M_\alpha}(S_i)$ for each $b_i \in \mathcal{B}$.

Afterwards, we assign sets $S_1, \dots, S_n$ to buyer $b_1, \dots, b_n$ increasing buyer b_i 's payment relative to the payment in the one-sided mechanism by the sum of the samples v'_s for $s \in S_i$ and pay each seller $s \in \hat{S}$ whose item has been sold the respective sample v'_s .

The construction given in our mechanism is stated in the value-oracle model, and direct computation of the adjusted valuations would be inefficient. We provide a discussion on how to implement the mechanism efficiently in Section 5.3. For purposes of our analysis, we assume that the one-sided mechanism M_α always assigns each buyer an inclusion-minimal set of items giving the according buyer at least the same utility (this can, e.g., be ensured by employing a simple type of tie-breaking which favors small sets over larger ones).

5.2 Analysis of SURPLUS MECHANISM

We start by restating Theorem 4, then we prove it by arguing in Lemma 4 that the two-sided mechanisms inherit the IR and DSIC properties from the one-sided mechanism and finally showing the welfare approximation in Lemma 5.

THEOREM 4 (BLACK BOX I). *Denote by α the approximation guarantee of any one-sided IR, DSIC offline/online mechanism for maximizing social welfare for XOS valuations. We give a two-sided mechanism*

Algorithm 1: SURPLUS MECHANISM

Set $\hat{S} = \emptyset, \hat{E} = \emptyset$
for all $s \in S$ **do**
 Propose to s a price of v'_s
 if s accepts **then**
 Set $\hat{S} = \hat{S} \cup \{s\}$
for all $(b, S) \in B \times 2^{\hat{S}}$ **do**
 $\hat{v}_b(S) = \sum_{s \in \hat{S}} (a_{b,\hat{S}}(s) - v'_s)$,
 where $\hat{S} = \operatorname{argmax}_{S^* \subseteq S \cap \hat{S}} \{v_b(S^*) - \sum_{s \in S^*} v'_s\}$
 $\hat{E} = \hat{E} \cup \{(b, S)\}$
Let A be the assignment on $\hat{G}$ induced by running M_α on the hypergraph $\hat{G} = (\mathcal{B} \cup \hat{S}, \hat{E}, \hat{v}_B)$, with hyperedge weights $\hat{v}_B$, presenting buyers to M_α according to its input requirements (e.g., offline or in random order)
for all $(b, S) \in A$ **do**
 b pays price $p_{(b,S)} = p_b^{M_\alpha}(S) + \sum_{s \in S} v'_s$, where $p_b^{M_\alpha}(S)$ is the price charged to b by M_α
 b gets assigned the items in S
 for each $s \in S$ **do**
 s receives a payment of v'_s

for XOS buyers and unit-supply sellers that is IR, DSIC, BB, uses a single sample from each seller and provides a $\max\{2\alpha, 3\}$ -approximation to the optimal social welfare. The two-sided mechanism inherits the offline/online properties of the one-sided mechanism on the buyer side and is offline on the seller side.

We start by establishing the individual rationality, truthfulness, and budget balance properties of the two-sided mechanism claimed in Theorem 4.

LEMMA 4. *Given that an IR and DSIC one-sided mechanism M_α is used, the two-sided SURPLUS MECHANISM is DSIC, IR, and WBB.*

PROOF. Let's fix any realization of the valuations. We start by showing truthfulness. Fix a seller $s \in S$: the only interaction s has with the algorithm is by accepting or rejecting the posted price v'_s , which is independent of v_s , in exchange for her item, which is clearly truthful and individually rational.

Fixing a buyer b : the algorithm will ask about her valuation and modify it to

$$\hat{v}_b(S) = \sum_{s \in \hat{S}} (a_{b,\hat{S}}(s) - v'_s), \text{ where } \hat{S} = \operatorname{argmax}_{S^* \subseteq S \cap \hat{S}} \left\{ v_b(S^*) - \sum_{s \in S^*} v'_s \right\},$$

reflecting that any item s will cost at least v'_s . This is done because M_α only involves items, not sellers, so the extra v'_s is charged afterwards.

$$\begin{aligned} & \operatorname{argmax}_{s \in S'} \{v_b(S') - \sum_{s \in S'} v'_s - p_b^{M_\alpha}(S') \mid v_b(S') \geq \sum_{s \in S'} v'_s + p_b^{M_\alpha}(S')\} \\ &= \operatorname{argmax}_{s \in S'} \{\hat{v}_b(S') - p_b^{M_\alpha}(S') \mid \hat{v}_b(S') \geq p_b^{M_\alpha}(S')\} \end{aligned}$$

where the argmax s are restricted to available S' and with the left hand side reflecting that M_α is truthful and the right hand side being exactly what buyer b is trying to maximize. Given that no trade is generated if for all sets $v_b(S') < \sum_{s \in S'} v'_s$, the mechanism is also individually rational. For any trade, the seller's price is v'_s

and the buyer's $v'_s + p_b^{M_\alpha}(s) \geq v'_s$, so the mechanism is budget balanced. $\square$

We now give the proof of the approximation ratio in Theorem 4.

LEMMA 5. *Assuming that an α -approximate one-sided mechanism M_α is used, the two-sided SURPLUS MECHANISM is $\max\{3, 2\alpha\}$ -approximate.*

PROOF. Recall that the adjusted buyer valuations in the graph $\hat{G}$ were defined as in Equation (1). Fix a pair (b, S) of buyer and set of items and a realization v_B of buyers' valuations. Then, in expectation over the sellers' valuations and samples, it holds that:

$$\mathbb{E}[v'_s \mid s \in \hat{S}] = \mathbb{E}[\max\{v_s, v'_s\}] \leq \mathbb{E}[2v_s], \quad (2)$$

where we used that $s \in \hat{S}$ if and only if $v_s \leq v'_s$. Note that for every pair (b, S') , and every realization of buyer and seller valuations and samples implies an according maximizing set in the definition of $\hat{v}_b(S')$, which we will denote as $\hat{S}(S')$. With this, we show

$$\begin{aligned} \mathbb{E}[\hat{v}_b(S)] &= \mathbb{E}\left[\sum_{s \in \hat{S}(S)} (a_{b,\hat{S}(S)}(s) - v'_s)\right] \\ &\geq \mathbb{E}\left[\sum_{s \in S \cap \hat{S}} (a_{b,S}(s) - v'_s)_+\right] = \sum_{s \in S \cap \hat{S}} \mathbb{E}[(a_{b,S}(s) - v'_s)_+] \\ &= \sum_{s \in S} \mathbb{E}[(a_{b,S}(s) - v'_s)_+ \mid s \in \hat{S}] \mathbb{P}(s \in \hat{S}) \\ &\geq \frac{1}{2} \sum_{s \in S} \mathbb{E}[a_{b,S}(s) - v'_s \mid s \in \hat{S}] \geq \frac{1}{2} \sum_{s \in S} \mathbb{E}[a_{b,S}(s) - 2v_s] \\ &= \frac{1}{2} \sum_{s \in S} \mathbb{E}[a_{b,S}(s)] - \sum_{s \in S} \mathbb{E}[v_s] = \frac{1}{2} v_b(S) - \sum_{s \in S} \mathbb{E}[v_s]. \end{aligned}$$

The first inequality follows from the fact that since $\hat{S}$ maximizes $v_b(S^*) - \sum_{s \in S^*} v'_s$, always choosing those $s \in S \cap \hat{S}$ for which $a_{b,S}(s) - v'_s \geq 0$ can only perform worse. Basic transformations and the fact that every $s \in S$ is included in $\hat{S}$ with probability at least $\frac{1}{2}$ result in the second inequality, at which point we simply plug in Inequality (2).

For the social welfare induced by a maximum-welfare assignment OPT in graph G , we have

$$SW_{OPT} = \sum_{(b,S') \in OPT} v_b(S') + \sum_{s \in S \setminus OPT} v_s.$$

Note that in our mechanism, each buyer will only be assigned a valuation-maximizing, inclusion-minimal set $\hat{S}$ from the definition of $\hat{v}_b$. Denote by OPT_1 an optimal assignment (hypermatching) in the original graph G , i.e., an optimal solution to the one-sided problem with original valuations v_b . Then, for the social welfare induced by our assignment A over $\hat{G}$, it holds

$$\begin{aligned} \mathbb{E}[SW_A] &= \mathbb{E}\left[\sum_{(b,\hat{S}) \in A} v_b(\hat{S}) + \sum_{s \in S \setminus A} v_s\right] \\ &\geq \mathbb{E}\left[\sum_{(b,\hat{S}) \in A} \hat{v}_b(\hat{S}) + \sum_{s \in S} v_s\right] \end{aligned}$$

$$\begin{aligned}
&\geq \mathbb{E} \left[\sum_{(b,S) \in OPT_1} \frac{1}{\alpha} \hat{v}_b(S) + \sum_{s \in S} v_s \right] \\
&\geq \mathbb{E} \left[\sum_{(b,S) \in OPT_1} \frac{1}{\alpha} \left(\frac{1}{2} v_b(S) - \sum_{s \in S} v_s \right) + \sum_{s \in S} v_s \right] \\
&= \mathbb{E} \left[\sum_{(b,S) \in OPT_1} \frac{1}{2\alpha} v_b(S) - \sum_{s \in OPT_1} \frac{1}{\alpha} v_s + \sum_{s \in S} v_s \right] \\
&\geq \mathbb{E} \left[\sum_{(b,S) \in OPT_1} \frac{1}{2\alpha} v_b(S) + \left(1 - \frac{1}{\alpha}\right) \sum_{s \in S} v_s \right] \\
&\geq \mathbb{E} \left[\sum_{(b,S) \in OPT} \frac{1}{2\alpha} v_b(S) + \left(1 - \frac{1}{\alpha}\right) \sum_{s \in S} v_s \right].
\end{aligned}$$

Here, the first inequality stems from the fact that the $\hat{v}_b$ are already discounted by $\sum_{s \in \hat{S}} v'_s \geq v_s$. Since M_α is an α -approximation to the optimum in $\hat{G}$, also the second inequality also holds true. Finally, we plug in the lower bound to $\hat{v}_b$ proven above. The expectation here is taken over all realizations of buyer valuations, seller valuations, and samples of seller valuations. For $\alpha \geq \frac{3}{2}$, from the previous inequality, one gets

$$\mathbb{E}[SW_A] \geq \frac{1}{2\alpha} \mathbb{E}[OPT].$$

For any other $\alpha \in [1, \frac{3}{2}]$, we observe that M_α is also a $\frac{3}{2}$ approximation and then the same calculation gives that $\mathbb{E}[SW_A] \geq \frac{1}{3} \mathbb{E}[OPT]$. So, all in all, we get a $\max\{2\alpha, 3\}$ -approximation, as claimed. $\square$

5.3 Computational Aspects

As already mentioned earlier, the adjusted valuations $\hat{v}_{\mathcal{B}}$ are stated with an argmax over subsets of S , which—depending on the computational model—may not be an efficient operation.

An alternative, which works for GS, is, for all $S \subseteq \hat{S}$, to define $\bar{v}_b(S) = \left(\sum_{s \in S \cap \hat{S}} (a_{b,S}(s) - v'_s) \right)_+$. As before it can be shown that if the original valuations are GS, then the modified valuations are GS as well. A difference between $\bar{v}_{\mathcal{B}}$ and $\hat{v}_{\mathcal{B}}$ is that the former need not be monotone, but monotonicity is not required by poly-time algorithms for GS valuations (see, e.g., [36]). Moreover, the approximation guarantee of the one-sided mechanism run on $\bar{v}_{\mathcal{B}}$ also applies if the resulting assignment is evaluated with the original adjusted valuations $\hat{v}_{\mathcal{B}}$ and the benchmark is the optimal allocation under the original adjusted valuations $\hat{v}_{\mathcal{B}}$, which is the only property of the one-sided mechanism that we used in our proof above. This is clarified in the following proposition.

PROPOSITION 2. *Any inclusion-minimal assignment ALG made by the algorithm run with adjusted valuations $\bar{v}$ instead of $\hat{v}$ provides the same approximation to the optimum welfare.*

PROOF. Note that the $\bar{v}_b$ are in the class XOS, so M_α 's approximation guarantee holds up. Also, for any pair (b, S) with $b \in B$, $S \subseteq \hat{S}$: $\hat{v}_b(S) \geq \bar{v}_b(S)$, since $\hat{v}_b(S)$ results from picking an optimal subset from S .

Let S^{opt} be an inclusion-minimal, utility-preserving set of items assigned to buyer b , which is a subset of b 's assigned bundle S^b in

an optimal hypermatching $OPT(\hat{G})$ of graph $\hat{G}$ as defined in our algorithm, i.e.,

$$\bar{v}_b(S^{opt}) = \sum_{s \in S^{opt}} (a_{b,S^{opt}}(s) - v'_s),$$

$$\text{where } S^{opt} = \operatorname{argmax}_{S^* \subseteq S^b \cap \hat{S}} \left\{ v_b(S^*) - \sum_{s \in S^*} v'_s \right\}.$$

It holds by definition

$$\bar{v}(S^{opt}) = \sum_{s \in S^{opt}} (a_{b,S^{opt}}(s) - v'_s)$$

and therefore, the weight of an optimal assignment is the same for both $\bar{v}$ and $\hat{v}$. Combining these facts, we get for the weight of the hypermatching $A_{\bar{v}}$ returned by the algorithm when using the valuations $\bar{v}_{\mathcal{B}}$:

$$\sum_{(b,S) \in A_{\bar{v}}} \hat{v}_b(S) \geq \sum_{(b,S) \in A_{\bar{v}}} \bar{v}_b(S) \geq c \cdot OPT(\bar{G}) = c \cdot OPT(\hat{G})$$

and this is what we needed in the proof of the approximation ratio. $\square$

In addition, note that while our mechanism is stated for the case of value queries, it can be formulated also for demand queries—as required by poly-time mechanisms for XOS valuations such as [1, 2, 20]—by adjusting the prices proposed instead of the buyer valuations. Here, it is simply necessary to increase the M_α -prices of any $S \subseteq \hat{S} \cap \hat{S}$ by $\sum_{s \in S} v'_s$. This does not reflect the *capping* of our adjusted valuations to a minimum contribution of 0 for each seller, but—having the same effect—buyers will never demand the according items.

Finally, observe that if we are allowed to use demand queries, then we can also efficiently implement value queries (as required by some demand-oracle algorithms) to the adjusted valuations $\hat{v}_{\mathcal{B}}$: first issue a demand query to find the set S^* in the argmax, and then issue a value query to obtain the value $v_b(S^*)$ of the corresponding set.

6 COMPUTATIONALLY EFFICIENT DOUBLE AUCTIONS

We now move on to less general settings with unit-demand buyers and unit-supply sellers with *identical items*. We present a general technique for turning truthful one-sided mechanisms into truthful two-sided mechanisms while enforcing the strong budget balance property (in contrast to just weak budget balance from the previous section).

The one-sided mechanisms we consider are for so-called binary single-parameter problems. In such a problem, an agent i can either win or lose, and has a value v_i for winning. The set of agents that can simultaneously win is given by a set system $\mathcal{I}$. The social welfare of a feasible set $X \in \mathcal{I}$ is simply the sum $\sum_{i \in X} v_i$ of the winning agents' valuations.

Given two set systems $\mathcal{I}$ and $\mathcal{I}'$ on a ground set U , we define its intersection to be the set system that contains all sets $X \subseteq U$ such that $X \in \mathcal{I}$ and $X \in \mathcal{I}'$.

THEOREM 5 (BLACK BOX II). *Denote by α the approximation guarantee of any one-sided IR, DSIC offline/online mechanisms for the*

intersection of a downward closed set system $\mathcal{I}$ with a uniform matroid that may or may not use information on the priors. We give a two-sided mechanism for unit-demand buyers, unit-supply sellers with identical goods, and constraints $\mathcal{I}$ on the buyers, that is IR, DSIC, SBB, has the same information requirements on the buyer side as the one-sided mechanism and uses a single sample from each seller, and yields a $(1 + 1/(2 - \sqrt{3}) \cdot \alpha) \approx (1 + 3.73 \cdot \alpha)$ approximation to the optimal social welfare. The mechanism inherits the online/offline properties of the one-sided mechanism on the buyer side and it is online random order on the seller side.

Instead of directly proving [Theorem 5](#), we present the algorithm for the following special case, and then argue how to generalize it.

THEOREM 6. *Let $k = \min\{n, m\}$. There is an IR, DSIC, SBB, $1 + \frac{1}{2-\sqrt{3}}(1 + O(1/\sqrt{k}))$ -approximate single-sample mechanism for unconstrained double auctions that approaches the buyers in online fixed order and the sellers in online random order.*

We give the explicit mechanism RESERVE REHEARSAL for [Theorem 6](#) in [Algorithm 2](#). RESERVE REHEARSAL runs the one-sided REHEARSAL algorithm [3] to select the top $k = \min\{n, m\}$ buyers. Azar et al. [3] show that the combined value of the buyers that beat their price provides, in expectation, a $1 + O(\frac{1}{\sqrt{k}})$ approximation to the expected value of the k highest buyers.

Our twist to this mechanism is that we pair buyers b that would be selected by the REHEARSAL algorithm with a random seller s , offering them to trade at a price that is the max of the respective buyer's $p_{\min}$ and the respective seller's sample v'_s . This adds the needed component to take into account the valuations on the seller side of the market, and will serve as an insurance that any good trade we propose has a good chance of actually happening.

Algorithm 2: RESERVE REHEARSAL

Let P be the set of the $k - 2\sqrt{k}$ largest buyer samples, together with $2\sqrt{k}$ copies of the $(k - 2\sqrt{k})^{th}$ -largest buyer sample
 Let $p_{\min}$ be the smallest element of P
 Fix any order on the buyers (or assume buyers arrive online)
for each b_i , **in this order do**
 if $v_{b_i} > p_{\min}$ **then**
 Delete from P the highest value $p \in P$ such that $v_{b_i} > p$
 Pick uniformly at random $s \in \mathcal{S}$, delete s from $\mathcal{S}$ (or assume sellers arrive online)
 Propose a trade to (b_i, s) for the price of $\max\{v'_s, p_{\min}\}$
 if b_i and s both agree **then**
 Make the trade at this price

For the more general result in [Theorem 5](#) we run the given one-sided mechanism on the intersection of the given feasibility constraint $(\mathcal{B}, \mathcal{I}_{\mathcal{B}})$ with a m -uniform matroid. This gives a set of tentative buyers B' along with tentative buyer prices $p_{\mathcal{B}}$. We can then randomly match the tentative buyers to sellers, offering buyer-seller pairs (b, s) to trade at price $\max\{p_b, v'_s\}$.

We begin by showing that our two-sided mechanism inherits IR and DSIC from its one-sided counterpart, and that it is strongly budget balanced.

LEMMA 6. *The RESERVE REHEARSAL mechanism is IR, DSIC, and SBB.*

PROOF. In RESERVE REHEARSAL no money is ever received by the mechanism itself. The only exchange of money happens between buyer-seller pairs (b_i, s) that also exchange an item. The mechanism is therefore strongly budget balanced.

It is also clear that the mechanism is individually rational as buyers and sellers would only accept trades at prices that are lower resp. higher than their respective valuations.

Furthermore, the mechanism is truthful for agents on both sides of the market. Each buyer is presented with a trading opportunity once; and the price depends only on the samples and the valuations of previously considered buyers. She can only accept or reject, but never influence it—and will therefore not profit from reporting a lower or higher value. Sellers, on the other hand, are also guaranteed to be considered only once. They, too, have no means of influencing the price they are presented with, and can only accept or reject. $\square$

It remains to show the claimed approximation guarantee. Just as in the case of bilateral trade, we will do the bulk of the work for fixed buyer valuations.

In what follows, we use $M = (x, p)$ to refer to the one-sided version of REHEARSAL, and $M' = (x', p')$ to refer to the two-sided version. We use B' to denote the set of tentative buyers chosen by M , we use B'_+ to denote the set of buyers that end up with an item in M' , and we use S_+ to denote the set of sellers that keep their item in M' . The expected social welfare achieved by M' is:

$$\mathbb{E}[\text{SW}(x'(v_{\mathcal{B}}, v_{\mathcal{S}}))] = \mathbb{E}\left[\sum_{b \in B'_+} v_b + \sum_{s \in S_+} v_s\right].$$

The key bit in our proof is the following lemma, which relates the performance of the one-sided mechanism to that of the two-sided mechanism.

LEMMA 7. *Let B' denote the set of tentative buyers chosen by the one sided mechanism M , let B'_+ denote the set of buyers that trade in the two-sided mechanism M' , and let S_+ denote the set of sellers that keep their item in the two-sided mechanism M' . Then,*

$$\mathbb{E}\left[\sum_{b \in B'_+} v_b + \sum_{s \in S_+} v_s\right] \geq (2 - \sqrt{3}) \cdot \mathbb{E}\left[\sum_{b \in B'} v_b\right].$$

PROOF. In order to prove the lemma we will show that for any fixed buyer valuations $v_{\mathcal{B}}$, buyer samples $v'_{\mathcal{B}}$, and corresponding set of tentative buyers B' , in expectation over the seller valuations $v_{\mathcal{S}}$, the seller samples $v'_{\mathcal{S}}$, and the randomness in the pairing of buyers and sellers,

$$\mathbb{E}\left[\sum_{b \in B'_+} v_b + \sum_{s \in S_+} v_s \mid B', v_{\mathcal{B}}, v'_{\mathcal{B}}\right] \geq (2 - \sqrt{3}) \cdot \sum_{b \in B'} v_b.$$

The actual claim then follows by taking expectation over buyer valuations $v_{\mathcal{B}}$, buyer samples $v'_{\mathcal{B}}$, and the corresponding set of tentative buyers B' .

In order to do the analysis let's fix any buyer $b \in B'$ with her value v_b and tentative payment p_b . Let's denote by s the random

seller associated to b , and by v_s and v'_s two independent samples from that seller's value distribution. Note that from buyer b 's perspective seller s is just a uniform random seller from $\mathcal{S}$.

The contribution of the couple (b, s) to the social welfare is v_b if there is a trade and v_s otherwise, and there is a trade when $v_b \geq \max\{p_b, v'_s\} \geq v_s$.

To analyze this contribution, we fix some constant $t \geq 0$ which will be set later and we denote with $F(x) = \mathbb{P}(v_s \leq x)$ the probability that a sample from F_s of a seller s chosen uniformly at random from $\mathcal{S}$ is at most x . For any t we can do a case distinction based on whether $v_b \geq t$ or $v_b < t$. If $v_b \geq t$, then the probability to have a trade is at least

$$\begin{aligned} & \mathbb{P}(\{v'_s \leq v_b\} \cap \{v_s \leq \max\{p_b, v'_s\}\} \mid t \leq v_b) \\ & \geq \mathbb{P}(\{v'_s \leq t\} \cap \{v_s \leq v'_s\}) \\ & \geq \mathbb{P}(\{v'_s \leq t\} \cap \{v_s \leq v'_s\} \cap \{v_s \leq t\}) \\ & \geq \mathbb{P}(v_s \leq v'_s \mid \{v'_s \leq t\} \cap \{v_s \leq t\}) \mathbb{P}(\{v_s \leq t\} \cap \{v'_s \leq t\}) \\ & \geq \frac{1}{2} F(t)^2. \end{aligned}$$

So the expected contribution is at least $\frac{1}{2} F(t)^2 v_b$. For the case $v_b < t$ we use the expected value of the seller as a lower bound to obtain that the expected contribution is at least

$$\begin{aligned} \mathbb{E}[v_s \mid t > v_b] &= \mathbb{E}[v_s] = \mathbb{E}[v'_s] \geq \mathbb{E}[v'_s \mid v'_s \geq v_b] \cdot \mathbb{P}(v'_s \geq v_b) \\ &\geq (1 - F(v_b)) v_b \geq \left(1 - F\left(\frac{t+v_b}{2}\right)\right) v_b. \end{aligned}$$

So the expected contribution of buyer b and her random partner s to the social welfare is at least $v_b \cdot LB(t)$, where

$$LB(t) \geq \begin{cases} \frac{1}{2} F(t)^2 & \text{if } t \leq v_b \\ 1 - F\left(\frac{v_b+t}{2}\right) & \text{if } t > v_b \end{cases}$$

We now set the arbitrary parameter $t \geq 0$ to lowerbound that term. Consider $t^* = \min\{t \mid F(t) \geq \sqrt{3} - 1\}$, which is the (possibly jumping) point of function F for which $F(t^*) \geq \sqrt{3} - 1$, and $F(t') < \sqrt{3} - 1$ for all $t' < t^*$. Using this t^* we can lower bound the expected contribution of B_i and s as follows:

$$\begin{aligned} v_b \cdot LB(t) &\geq v_b \cdot \min\left\{\frac{1}{2} F(t^*)^2, 1 - F\left(\frac{v_b+t^*}{2}\right)\right\} \\ &\geq v_b \cdot \min\left\{\frac{1}{2} 2(2 - \sqrt{3}), 2 - \sqrt{3}\right\} \\ &= v_b \cdot \min\{2 - \sqrt{3}, 2 - \sqrt{3}\} = v_b \cdot (2 - \sqrt{3}). \end{aligned}$$

One concise way to express the progress so far is:

$$\mathbb{E}\left[\sum_{a \in (B'_+ \cup S_+) \cap \{b, s\}} v_a \mid B', v_B, v'_B\right] \geq (2 - \sqrt{3}) \cdot v_b, \quad \forall b \in B'$$

where the randomness is over the choice of the random seller s and her values v_s and v'_s . The above holds for all buyers b , so we can sum up for all buyers in B' , then use linearity of expectation and the fact that $|B'| \leq m$ to obtain

$$\begin{aligned} \sum_{b \in B'} (2 - \sqrt{3}) \cdot v_b &\leq \sum_{b \in B'} \mathbb{E}\left[\sum_{a \in (B'_+ \cup S_+) \cap \{b, s\}} v_a \mid B', v_B, v'_B\right] \\ &\leq \mathbb{E}\left[\sum_{b \in B'_+} v_b + \sum_{s \in S_+} v_s \mid B', v_B, v'_B\right], \end{aligned}$$

as claimed. $\square$

We are now ready to prove the performance guarantee of the two-sided mechanism M' . This, together with [Lemma 6](#), concludes the proof of [Theorem 6](#).

LEMMA 8. Fix n and m and let α denote the approximation guarantee of the REHEARSAL algorithm. The RESERVE REHEARSAL mechanism yields in expectation a

$$\left(1 + \frac{1}{2 - \sqrt{3}} \cdot \alpha\right) \approx (1 + 3.73 \cdot \alpha)$$

approximation to the expected optimal social welfare.

PROOF. Recall that we use B'_+ to denote the set of buyers that actually do get an item in our two-sided mechanism M' , and that we use S_+ to denote those sellers that do not make any trade and keep their item. With this notation the expected social welfare achieved by our two-sided mechanism M' is

$$\mathbb{E}[\text{SW}(x'(v_B, v_S))] = \mathbb{E}\left[\sum_{b \in B'_+} v_b + \sum_{s \in S_+} v_s\right].$$

For a given set of valuations v_B of the buyers, denote by $OPT_k(v_B)$ the set of buyers with the k highest values. We can upper bound the expected optimal social welfare by the optimal solution for the buyers plus all seller values

$$\mathbb{E}[\text{SW}(OPT(v_B, v_S))] \leq \mathbb{E}\left[\sum_{b \in OPT_k(v_B)} v_b\right] + \mathbb{E}\left[\sum_{s \in S} v_s\right]. \quad (3)$$

Recall that the one-sided mechanism M computes a set B' of buyers whose accumulated expected values are at least $\frac{1}{\alpha}$ times the expected one-sided optimum. Hence, for the considered buyers B' ,

$$\frac{1}{\alpha} \cdot \mathbb{E}\left[\sum_{b \in OPT_k(v_B)} v_b\right] \leq \mathbb{E}\left[\sum_{b \in B'} v_b\right].$$

By combining the above inequality with our upper bound on the expected optimal social welfare in [Inequality \(3\)](#), we obtain

$$\mathbb{E}[\text{SW}(OPT(v_B, v_S))] \leq \alpha \cdot \mathbb{E}\left[\sum_{b \in B'} v_b\right] + \mathbb{E}\left[\sum_{s \in S} v_s\right]. \quad (4)$$

First consider the second term on the right hand side of [Inequality \(4\)](#). In our two-sided mechanism M' sellers trade only if they are matched to a buyer with higher valuation. Therefore, we can replace as follows:

$$\mathbb{E}[\text{SW}(OPT(v_B, v_S))] \leq \alpha \cdot \mathbb{E}\left[\sum_{b \in B'} v_b\right] + \mathbb{E}\left[\sum_{s \in S_+} v_s + \sum_{b \in B'_+} v_b\right].$$

Now consider the first term on the right hand side of [Inequality \(4\)](#). Our two-sided mechanism M' does not make trades for each buyer in B' . Despite the fact that generally $B' \neq B'_+$, as we show in [Lemma 7](#), in expectation, $(2 - \sqrt{3}) \cdot \mathbb{E}\left[\sum_{b \in B'} v_b\right]$ is a lower bound on our mechanism's social welfare. Using this we obtain

$$\begin{aligned} \mathbb{E}[\text{SW}(OPT(v_B, v_S))] &\leq \frac{\alpha}{2 - \sqrt{3}} \cdot \mathbb{E}\left[\sum_{b \in B'_+} v_b + \sum_{s \in S_+} v_s\right] + \mathbb{E}\left[\sum_{s \in S_+} v_s + \sum_{b \in B'_+} v_b\right]. \end{aligned}$$

All in all, we get:

$$\mathbb{E}[\text{SW}(OPT(v_B, v_S))] \leq \left(1 + \frac{\alpha}{2 - \sqrt{3}}\right) \cdot \mathbb{E}\left[\sum_{b \in B'_+} v_b + \sum_{s \in S_+} v_s\right],$$

as claimed. $\square$

We conclude by describing how to extend this argument to prove the more general result in [Theorem 5](#).

SKETCH OF THE PROOF OF THEOREM 5. The proof that the general mechanism described above is IR, DSIC, and SBB is basically identical to that of [Lemma 6](#). The key is that truthfulness of the one-sided mechanism ensures that tentative buyers want an opportunity to trade, and cannot manipulate the price they face for this opportunity. For the performance analysis we claim that [Lemma 8](#) applies more generally with α being the approximation guarantee of the one-sided mechanism. In fact, the only change to the above proof that is required for this generalization is to redefine OPT_k as the optimal one-sided solution containing at most $k = \min\{n, m\}$ buyers. $\square$

7 CONCLUSION AND FURTHER DIRECTIONS

We have initiated the study of simple mechanisms for two-sided markets that use only a minimum amount of information from the priors, i.e., just a single sample from the sellers' distribution. While without prior information, there is generally no hope for any meaningful approximation, our mechanisms are approximately efficient up to small constant factors.

This line of research is of specific relevance for two-sided markets since efficient mechanisms with the desired requirements are only possible in the Bayesian setting and, moreover, the optimal Bayesian mechanism is complicated and known only for restricted cases. Our results are very general, and in several cases even improve on the best known approximation guarantee with perfect knowledge of the distributions. Our results can be extended further by identifying new efficient mechanisms for the one-sided versions of the problems. Finally, we leave open the problem of devising single sample mechanisms for the challenging problem of optimizing the *gain from trade* in two-sided markets.

A OMITTED PROOFS

PROOF OF THEOREM 1. We concentrate on bilateral trade only, since for richer two-sided problems, we can just assume the social welfare to be dominated by a single buyer-seller pair. Let α be any approximation factor larger than one and assume for contradiction the existence of an α -approximative mechanism M as described above. Call $\mathbb{P}_{v_b, v_s}$ the probability of trade for any reported v_b, v_s following M .

We start off by fixing some reported buyer valuation v_b , and argue that if the seller reports a value $v_s^* = \frac{v_b}{\alpha^2}$, mechanism M must trade the item from s to b with probability $\mathbb{P}_{v_b, v_s^*}$ at least $\frac{1}{1+\alpha}$. In fact, since M is an α -approximation, we have that the expected social welfare has to be at least a $1/\alpha$ factor of v_b :

$$\mathbb{P}_{v_b, v_s^*} \cdot v_b + \left(1 - \mathbb{P}_{v_b, v_s^*}\right) \cdot v_s^* \geq \frac{1}{\alpha} v_b$$

which leads to $\mathbb{P}_{v_b, v_s^*} \geq \frac{1}{\alpha+1}$. We will employ this bound on $\mathbb{P}_{v_b, v_s^*}$ to infer a lower bound for the expected price of trade of the seller p_s in terms of v_b and α . On one hand, when the seller reports v_s^* with her true valuation being $v_s \leq v_s^*$, she gains at least

$$g_s^* = \mathbb{P}_{v_b, v_s^*} \cdot \mathbb{E}[p_s(v_b, v_s^*) | \text{trade}] + (1 - \mathbb{P}_{v_b, v_s^*}) \cdot v_s$$

$$\geq \frac{1}{\alpha+1} v_s^* + \left(1 - \frac{1}{\alpha+1}\right) v_s \geq \frac{v_b}{(\alpha+1)\alpha^2},$$

where $p_s(v_b, v_s) \geq v_s$ by IR when a trade occurs. On the other hand, since M is truthful, reporting v_s gives the seller at least the same gain, i.e., $g_s \geq g_s^*$, where

$$g_s = \mathbb{P}_{v_b, v_s} \cdot \mathbb{E}[p_s(v_b, v_s) | \text{trade}] + (1 - \mathbb{P}_{v_b, v_s}) \cdot v_s \leq \mathbb{E}[p_s(v_b, v_s) | \text{trade}].$$

Put together, we get for every fixed v_b and $v_s \leq v_b/\alpha^2$,

$$\mathbb{E}[p_s(v_b, v_s) | \text{trade}] \geq g_s \geq g_s^* \geq \frac{v_b}{(\alpha+1)\alpha^2}. \quad (5)$$

Now we change perspective: we fix any v_s and consider any two values $v_b > v_b'$ for the buyer. As a first step in this new direction we show that the probability of a trade is a non-decreasing function of the reported buyer valuations. By truthfulness, assuming that the true value is v_b we obtain

$$(v_b - \mathbb{E}[p_b(v_b, v_s) | \text{trade}]) \mathbb{P}_{v_b, v_s} \geq (v_b - \mathbb{E}[p_b(v_b', v_s) | \text{trade}]) \mathbb{P}_{v_b', v_s},$$

similarly, if the true value is v_b'

$$(v_b' - \mathbb{E}[p_b(v_b', v_s) | \text{trade}]) \mathbb{P}_{v_b', v_s} \geq (v_b' - \mathbb{E}[p_b(v_b, v_s) | \text{trade}]) \mathbb{P}_{v_b, v_s}.$$

Adding both inequalities, we have that $\mathbb{P}_{v_b, v_s} \geq \mathbb{P}_{v_b', v_s}$. This implies that, for any fixed v_s , there is a fixed value $\mathbb{P}_{v_s}$ such that $\lim_{v_b \rightarrow \infty} \mathbb{P}_{v_b, v_s} = \mathbb{P}_{v_s} \leq 1$.

At the same time, we have that $\mathbb{E}[p_b(v_b, v_s) | \text{trade}]$ is also non-decreasing in v_b , otherwise reporting a higher value would be a beneficial deviation for the buyer, increasing the probability of a trade while keeping the expected price paid at most as large.

For any v_s , and any $\epsilon > 0$ arbitrarily close to 0, let v_b^{ϵ, v_s} denote a fixed buyer bid v_b such that the trading probability $\mathbb{P}_{v_b, v_s}$ for (v_b, v_s) is at least $\mathbb{P}_{v_s} - \epsilon$. Then, for a buyer value of $v_b \geq v_b^{\epsilon, v_s}$, truthfulness yields

$$\begin{aligned} & \mathbb{P}_{v_s} (v_b - \mathbb{E}[p_b(v_b, v_s) | \text{trade}]) \\ & \geq (\mathbb{P}_{v_s} - \epsilon) (v_b - \mathbb{E}[p_b(v_b^{\epsilon, v_s}, v_s) | \text{trade}]). \end{aligned}$$

With some simple algebra, we get that

$$\begin{aligned} & \left(1 - \left(\frac{\mathbb{P}_{v_s} - \epsilon}{\mathbb{P}_{v_s}}\right)\right) v_b \\ & \geq \mathbb{E}[p_b(v_b, v_s) | \text{trade}] - \left(\frac{\mathbb{P}_{v_s} - \epsilon}{\mathbb{P}_{v_s}}\right) \mathbb{E}[p_b(v_b^{\epsilon, v_s}, v_s) | \text{trade}] \\ & \geq \mathbb{E}[p_b(v_b, v_s) | \text{trade}] - \mathbb{E}[p_b(v_b^{\epsilon, v_s}, v_s) | \text{trade}] \end{aligned} \quad (6)$$

Having shown these inequalities about the expected trading prices for the buyer and seller, we need to combine both. To this end, consider any input (v_b, v_s) to mechanism M with $v_b \gg v_s$, i.e., $v_b \geq \max\{\alpha^2 v_s, v_b^{\epsilon, v_s}\}$. Weak budget-balance together with individual rationality implies, using [Inequality \(5\)](#) and rearranging the terms in [Inequality \(6\)](#):

$$\begin{aligned} & \left(1 - \frac{\mathbb{P}_{v_s} - \epsilon}{\mathbb{P}_{v_s}}\right) v_b - \mathbb{E}[p_b(v_b^{\epsilon, v_s}, v_s) | \text{trade}] \geq \mathbb{E}[p_b(v_b, v_s) | \text{trade}] \\ & \geq \mathbb{E}[p_s(v_b, v_s) | \text{trade}] \geq \frac{v_b}{(\alpha+1)\alpha^2}. \end{aligned}$$

Reordering this, we have that the expected price charged from the buyer when she reports v_b^ϵ is at least some factor times v_b , for any large v_b . More concretely, it holds

$$\mathbb{E} \left[p_b(v_b^{\epsilon, v_s}, v_s) | \text{trade} \right] \geq \left(\frac{1}{(\alpha + 1)\alpha^2} - \left(1 - \frac{\mathbb{P}_{v_s} - \epsilon}{\mathbb{P}_{v_s}} \right) \right) v_b.$$

However, for ϵ chosen small enough, the right side is clearly approaching infinity when v_b increases, while v_b^{ϵ, v_s} remains constant since it depends only on v_s but not v_b . We have therefore shown that $\mathbb{E} \left[p_b(v_b^{\epsilon, v_s}, v_s) | \text{trade} \right] > v_b^{\epsilon, v_s}$ for appropriate choice of ϵ and v_b , a contradiction to M being IR. $\square$

REFERENCES

- [1] S. Assadi, T. Kesselheim, and S. Singla. 2021. Improved Truthful Mechanisms for Subadditive Combinatorial Auctions: Breaking the Logarithmic Barrier. In *Proc. 32nd ACM-SIAM SODA*. 653–661. <https://doi.org/10.1137/1.9781611976465.40>
- [2] S. Assadi and S. Singla. 2019. Improved Truthful Mechanisms for Combinatorial Auctions with Submodular Bidders. In *Proc. 60th IEEE FOCS*. 233–248. <https://doi.org/10.1109/FOCS.2019.00024>
- [3] P. D. Azar, R. Kleinberg, and S. Matthew Weinberg. 2014. Prophet Inequalities with Limited Information. In *Proc. 25th ACM-SIAM SODA*. 1358–1377. <https://doi.org/10.1137/1.9781611973402.100>
- [4] M. Babaioff, Y. Cai, Y. A. Gonczarowski, and M. Zhao. 2018. The Best of Both Worlds: Asymptotically Efficient Mechanisms with a Guarantee on the Expected Gains-From-Trade. In *Proc. 19th ACM EC*. 373. <https://doi.org/10.1145/3219166.3219203>
- [5] M. Babaioff and N. Nisan. 2004. Concurrent Auctions Across the Supply Chain. *J. Artif. Intell. Res.* 21 (2004), 595–629. <https://doi.org/10.1613/jair.1316>
- [6] M. Babaioff and W. E. Walsh. 2005. Incentive-Compatible, Budget-Balanced, yet Highly Efficient Auctions for Supply Chain Formation. *Decis. Support Sys.* 39, 1 (2005), 123–149. <https://doi.org/10.1016/j.dss.2004.08.008>
- [7] S. R. Balseiro, V. S. Mirrokni, R. Paes Leme, and S. Zuo. 2019. Dynamic Double Auctions: Towards First Best. In *Proc. 30th ACM-SIAM SODA*. 157–172. <https://doi.org/10.1137/1.9781611975482.11>
- [8] L. Blumrosen and S. Dobzinski. 2014. Reallocation Mechanisms. In *Proc. 15th ACM EC*. 617. <https://doi.org/10.1145/2600057.2602843>
- [9] L. Blumrosen and S. Dobzinski. 2016. (Almost) Efficient Mechanisms for Bilateral Trading. *CoRR* (2016). <http://arxiv.org/abs/1604.04876>
- [10] L. Blumrosen and Y. Mizrahi. 2016. Approximating Gains-from-Trade in Bilateral Trading. In *Proc. 12th WINE*. 400–413. https://doi.org/10.1007/978-3-662-54110-4_28
- [11] J. Brustle, Y. Cai, F. Wu, and M. Zhao. 2017. Approximating Gains from Trade in Two-sided Markets via Simple Mechanisms. In *Proc. 20th ACM EC*. 589–590. <https://doi.org/10.1145/3033274.3085148>
- [12] E. H. Clarke. 1971. Multipart pricing of public goods. *Public Choice* 11, 1 (1971), 17–33. <https://doi.org/10.1007/BF01726210>
- [13] R. Cole and T. Roughgarden. 2014. The Sample Complexity of Revenue Maximization. In *Proc. 46th ACM-SIGACT STOC*. 243–252. <https://doi.org/10.1145/2591796.2591867>
- [14] R. Colini-Baldeschi, B. de Keijzer, S. Leonardi, and S. Turchetta. 2016. Approximately Efficient Double Auctions with Strong Budget Balance. In *Proc. 27th ACM-SIAM SODA*. 1424–1443. <https://doi.org/10.1137/1.9781611974331.ch98>
- [15] R. Colini-Baldeschi, P. W. Goldberg, B. de Keijzer, S. Leonardi, T. Roughgarden, and S. Turchetta. 2017. Approximately Efficient Two-Sided Combinatorial Auctions. In *Proc. 18th ACM EC*. 591–608. <https://doi.org/10.1145/3033274.3085128>
- [16] R. Colini-Baldeschi, P. W. Goldberg, B. de Keijzer, S. Leonardi, and S. Turchetta. 2017. Fixed Price Approximability of the Optimal Gain from Trade. In *Proc. 13th WINE*. 146–160. https://doi.org/10.1007/978-3-319-71924-5_11
- [17] J. R. Correa, A. Cristi, B. Epstein, and J. A. Soto. 2020. The Two-Sided Game of Googol and Sample-Based Prophet Inequalities. In *Proc. 31st ACM-SIAM SODA*. 2066–2081. <https://doi.org/10.1137/1.9781611975994.127>
- [18] J. R. Correa, P. Dütting, F. A. Fischer, and K. Schewior. 2019. Prophet Inequalities for I.I.D. Random Variables from an Unknown Distribution. In *Proc. 20th ACM EC*. 3–17. <https://doi.org/10.1145/3328526.3329627>
- [19] P. Dhangwatnotai, T. Roughgarden, and Q. Yan. 2015. Revenue maximization with a single sample. *Games and Econ. Behav.* 91 (2015), 318–333. <https://doi.org/10.1016/j.geb.2014.03.011>
- [20] S. Dobzinski. 2016. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. In *Proc. 48th ACM-SIGACT STOC*. 940–948. <https://doi.org/10.1145/2897518.2897569>
- [21] P. Dütting, T. Roughgarden, and I. Talgam-Cohen. 2017. Modularity and Greed in Double Auctions. *Games Econ. Behav.* 105 (2017), 59–83. <https://doi.org/10.1016/j.geb.2017.06.008>
- [22] Y. Giannakopoulos, E. Koutsoupias, and P. Lazos. 2017. Online Market Intermediation. In *Proc. 44th ICALP*. 47:1–47:14. <https://doi.org/10.4230/LIPIcs.ICALP.2017.47>
- [23] K. Goldner, M. Babaioff, and Y. A. Gonczarowski. 2020. Bulow-Klemperer-Style Results for Welfare Maximization in Two-Sided Markets. In *Proc. 31st ACM-SIAM SODA*. 2452–2471. <https://doi.org/10.1137/1.9781611975994.150>
- [24] R. Gomes and V. S. Mirrokni. 2014. Optimal Revenue-Sharing Double Auctions with Applications to Ad Exchanges. In *Proc. 23rd WWW*. 19–28. <https://doi.org/10.1145/2566486.2568029>
- [25] T. Groves. 1973. Incentives in Teams. *Econometrica* 41, 4 (1973), 617–631. <https://doi.org/10.2307/1914085>
- [26] Z. Y. Kang and J. Vondrák. 2019. Fixed-Price Approximations to Optimal Efficiency in Bilateral Trade. *SSRN* (2019). <https://ssrn.com/abstract=3460336>
- [27] R. D. Kleinberg. 2005. A Multiple-Choice Secretary Algorithm with Applications to Online Auctions. In *Proc. 16th ACM-SIAM SODA*. 630–631.
- [28] E. Koutsoupias and P. Lazos. 2018. Online trading as a secretary problem. In *Proc. 11th SAGT*. 201–212. https://doi.org/10.1007/978-3-319-99660-8_18
- [29] R. P. McAfee. 1992. A Dominant Strategy Double Auction. *J. Econ. Theory* 56, 2 (1992), 434–450. [https://doi.org/10.1016/0022-0531\(92\)90091-U](https://doi.org/10.1016/0022-0531(92)90091-U)
- [30] R. P. McAfee. 2008. The Gains from Trade under Fixed Price Mechanisms. *Appl. Econ. Res. Bull.* 1, 1 (2008), 1–10.
- [31] J. Morgenstern and T. Roughgarden. 2015. On the Pseudo-Dimension of Nearly Optimal Auctions. In *Proc. NIPS*. 136–144.
- [32] A. Mu'alem and N. Nisan. 2008. Truthful Approximation Mechanisms for Restricted Combinatorial Auctions. *Games Econ. Behav.* 64, 2 (2008), 612–631. <https://doi.org/10.1016/j.geb.2007.12.009>
- [33] R. B. Myerson and M. A. Satterthwaite. 1983. Efficient Mechanisms for Bilateral Trading. *J. Econ. Theory* 29, 2 (1983), 265–281. [https://doi.org/10.1016/0022-0531\(83\)90048-0](https://doi.org/10.1016/0022-0531(83)90048-0)
- [34] R. Niazadeh, Y. Yuan, and R. D. Kleinberg. 2014. Simple and Near-Optimal Mechanisms for Market Intermediation. In *Proc. 10th WINE*. 386–399. https://doi.org/10.1007/978-3-319-13129-0_31
- [35] N. Nisan and I. Segal. 2006. The Communication Requirements of Efficient Allocations and Supporting Prices. *J. Econ. Theory* 129, 1 (2006), 192–224. <https://doi.org/10.1016/j.jet.2004.10.007>
- [36] R. Paes Leme. 2017. Gross Substitutability: An Algorithmic Survey. *Games Econ. Behav.* 106 (2017), 294–316. <https://doi.org/10.1016/j.geb.2017.10.016>
- [37] R. Reiffenhäuser. 2019. An Optimal Truthful Mechanism for the Online Weighted Bipartite Matching Problem. In *Proc. 30th ACM-SIAM SODA*. 1982–1993. <https://doi.org/10.1137/1.9781611975482.120>
- [38] A. Rubinstein, J. Z. Wang, and S. M. Weinberg. 2020. Optimal Single-Choice Prophet Inequalities from Samples. In *Proc. 11th ITCS*. 60:1–60:10. <https://doi.org/10.4230/LIPIcs.ITCS.2020.60>
- [39] A. Rustichini, M. A. Satterthwaite, and S. R. Williams. 1994. Convergence to efficiency in a simple market with incomplete information. *Econometrica* 62, 5 (1994), 1041–1063. <https://doi.org/10.2307/2951506>
- [40] M. A. Satterthwaite and S. R. Williams. 1989. The Rate of Convergence to Efficiency in the Buyer's Bid Double Auction as the Market Becomes Large. *Rev. Econ. Stud.* 56, 4 (1989), pp. 477–498. <https://doi.org/10.2307/2951506>
- [41] M. A. Satterthwaite and S. R. Williams. 2002. The optimality of a simple market mechanism. *Econometrica* 70, 5 (2002), 1841–1863. <https://www.jstor.org/stable/3082022>
- [42] E. Segal-Halevi, A. Hassidim, and Y. Aumann. 2016. SBBA: A Strongly-Budget-Balanced Double-Auction Mechanism. In *Proc. 9th SAGT*. 260–272. https://doi.org/10.1007/978-3-662-53354-3_21
- [43] W. Vickrey. 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. Finance* 16, 1 (1961), 8–37. <https://doi.org/10.2307/2977633>