Multi-Unit Bilateral Trade

Matthias Gerstgrasser,[1] Paul W. Goldberg,[2] Bart de Keijzer,[3] Philip Lazos,[4] Alexander Skopalik[5]
[1,2,4]Department of Computer Science, University of Oxford
[3]School of Computer Science and Electronic Engineering, University of Essex
[5]Department of Applied Mathematics, University of Twente
{matthias.gerstgrasser, paul.goldberg, filippos.lazos}@cs.ox.ac.uk
b.dekeijzer@essex.ac.uk
a.skopalik@utwente.nl

Abstract

We characterise the set of dominant strategy incentive compatible (DSIC), strongly budget balanced (SBB), and ex-post individually rational (IR) mechanisms for the multi-unit bilateral trade setting. In such a setting there is a single buyer and a single seller who holds a finite number k of identical items. The mechanism has to decide how many units of the item are transferred from the seller to the buyer and how much money is transferred from the buyer to the seller. We consider two classes of valuation functions for the buyer and seller: Valuations that are increasing in the number of units in possession, and the more specific class of valuations that are increasing and submodular.

Furthermore, we present some approximation results about the performance of certain such mechanisms, in terms of social welfare: For increasing submodular valuation functions, we show the existence of a deterministic 2-approximation mechanism and a randomised $e/(1-e)$ approximation mechanism, matching the best known bounds for the single-item setting.

1 Introduction

Auctions form one of the most studied applications of game theory and mechanism design. In an auction setting, a single seller or auctioneer runs a pre-determined procedure or mechanism (i.e., the auction) to sell one or more goods to the buyers, and the buyers then have to strategise on the way they interact with the auction mechanism. An auction setting is rather restrictive in that it involves a single seller that is monopolistic and is assumed to be non-strategic. While this is a sufficient assumption in some cases, there are many applications that are more complex: It is often realistic to assume that a seller expresses a valuation for the items in her possession and that a seller wants to maximise her profit. Such settings in which both buyers and sellers are considered as strategic agents are known as *two-sided markets*, whereas auction settings are often referred to as *one-sided markets*.

The present paper falls within the area of mechanism design for two-sided markets, where the focus is on designing satisfactory market platforms or intermediation mechanisms that enable trade between buyers and sellers. In general, the term "satisfactory" can be tailored to the specific market under consideration, but nonetheless, in economic theory various universal properties have been identified and agreed on as important. The following three are the most fundamental ones:

- *Incentive Compatibility ((DS)IC)*: It must be a dominant strategy for the agents (buyers and sellers) to behave truthfully, hence not "lie" about their valuations for the items in the market. This enables the market mechanism to make an informed decision about the trades to be made.
- *Individual Rationality (IR)*: It must not harm the utility of an agent to participate in the mechanism.
- *Strong Budget Balance (SBB)*: All monetary transfers that the mechanism executes are among participating agents only. That is, no money is injected into the market, and no money is burnt or transferred to any agent outside of the market.

This paper studies the capabilities of mechanisms that satisfy these three fundamental properties above for a very simple special case of a two-sided market. *Bilateral trade* is the most basic such setting comprising a buyer and a seller, together with a single item that may be sold, i.e., transferred from the seller to the buyer against a certain payment from the buyer to the seller. The bilateral trade setting is a classical one: It was studied in the seminal paper (Myerson and Satterthwaite 1983) and has been studied in detail in various other publications in the economics literature. Recent work in the Algorithmic Game Theory literature (Blumrosen and Dobzinski 2016; Blumrosen and Mizrahi 2016; Colini-Baldeschi et al. 2017b) has focused on the welfare properties of bilateral trade mechanisms. These works assume the existence of *prior distributions* over the valuations of the buyer and seller, that may be thought of as modelling an intermediary's beliefs about the buyer's and seller's values for the item.

The present paper studies a generalisation of the classical bilateral trade setting by allowing the seller to hold multiple units initially. These units are assumed to be of a single resource, so that both agents only express valuations in terms of how many units they have in possession. The final utility of an agent (buyer or seller) is then determined by her valuation and the payment she paid or received. We focus our study on characterising which mechanisms satisfy the the above three properties and which of these feasible mechanisms achieve a good social welfare (i.e., total utility of buyer and seller combined).

Due to its simplicity, our setting is fundamental to any strategic setting where items are to be redistributed or reallocated. Our characterisation efforts show that all feasible mechanisms must belong to a very restricted class, already for this very simple setting with one buyer, one seller, and a relatively simple valuation structure. The specific mechanisms we develop are very simple, and suitable for implementation with very little communication complexity.

Our Contribution. Our first main contribution is a full characterisation of the class of truthful, individually rational and strongly budget balanced mechanisms in this setting. We do this separately for two classes of valuation functions: submodular valuations and general non-decreasing valuations. Section 3 presents a high-level argument for the submodular case. A full and rigorous formal proof for both settings is given in the full version of this paper (Gerstgrasser et al. 2018). Essentially, for the general case, any mechanism that aims to be truthful, strongly budget balanced and individually rational can only allow the agents to trade a single quantity of items at a predetermined price. The trade then only occurs if both the seller and buyer agree to it. This leads to a very clean characterization and has the added benefit of giving a robust, simple to understand mechanism: the agents do not have to disclose their entire valuation to the mechanism, and only have to communicate whether they agree to trade one specific quantity at one specific price. For the submodular case, suitable mechanisms can be characterised as specifying a per-unit price, and repeatedly letting the buyer and seller trade an item at that price until one of them declines to continue.

Secondly, we give approximation mechanisms for the social welfare objective in the Bayesian setting in Section 4, for the case of submodular valuations. Theorem 4.1 presents a 2-approximate deterministic mechanism. For randomised mechanisms, we show a $e/(e-1)$-approximation in Theorem 4.2.

Related Literature. The first approximation result for bilateral trade was presented in (McAfee 2007), where for the single-item case the author proves that the optimal *gain from trade* can be 2-approximated by the *median mechanism*, which is a mechanism that sets the seller's median valuation as a fixed price for the item, and trade occurs if and only if p lies in between the buyer's and seller's valuation and the buyer's valuation exceeds p. The analysis in (McAfee 2007) is done under the assumption that the seller's median valuation does not exceed the median valuation of the buyer. The gain from trade is defined as the increase in social welfare as a result of trading the item. (Blumrosen and Dobzinski 2016) extended the analysis of this mechanism by showing that it also 2-approximates the social welfare without the latter assumption on the medians.

In (Blumrosen and Dobzinski 2016), the authors consider the classical bilateral trade setting (with a single item) and present various mechanisms for it that approximate the optimal social welfare. Their best mechanism achieves an approximation factor of $e/(e-1)$. As in the present paper, there are prior distributions on the traders' valuations, and the quantity being approximated is the expectation over the priors, of the optimal allocation of the item.

The weaker notion of Bayesian incentive compatibility is considered in (Blumrosen and Mizrahi 2016), where the authors propose a mechanism in which the seller offers a take-it-or-leave-it price to the buyer. They prove that this mechanism approximates the harder *gain from trade* objective within a factor of $1/e$ under a technical albeit often reasonable *MHR condition* on the buyer's distribution.

The class of DSIC, IR, and SBB mechanisms for bilateral trade was characterised in (Colini-Baldeschi et al. 2016) to be the class of *fixed price mechanisms*. In the present work, we characterise this set of mechanisms for the more general multi-unit bilateral trade setting, thereby extending their result. The gain from trade arising from such mechanisms was analysed in (Colini-Baldeschi et al. 2017a).

Various recent papers analyse more general two-sided markets, where there are multiple buyers and sellers, who hold possibly complex valuations over the items in the market. (Colini-Baldeschi et al. 2017b) analyse a more general scenario with multiple buyers, sellers, and multiple distinct items, and use the same feasibility requirements as ours (DSIC, IR, and SBB). (Segal-Halevi, Hassidim, and Aumann 2018b) have considered a similar setting but focus on *gains from trade (GFT)* (i.e., the increase in social welfare resulting from reallocation of the items) instead of welfare. They initially considered a multi-unit setting like ours (albeit with multiple buyers and sellers), and they extend their work in (Segal-Halevi, Hassidim, and Aumann 2018a) to allow multiple types of goods. They present a mechanism that approximates the optimal GFT asymptotically in large markets. (Balseiro et al. 2018) designs two-sided market mechanisms for one seller and multiple buyers with a temporal component, where valuations are correlated between buyers but independent across time steps. A good approximation (of factor $1/2$) of the social welfare using the more permissive notion of *Bayesian* Incentive Compatibility (BIC) was achieved by (Brustle et al. 2017). Their optimality benchmark is different from the one we consider as they compare their mechanism to the best possible BIC, IR, and SBB mechanism. A very recent work, (Babaioff et al. 2018), proposes mechanisms that achieve social welfare guarantees for both optimality benchmarks. (Feldman and Gonen 2018) considers optimizing the gains from trade in a two-sided market setting tailored to online advertising platforms, and the authors extend this idea further in (Feldman and Gonen 2016) by considering two-sided markets in an online setting.

The literature discussed so far aims to maximise welfare under some budget-balance constraints. An alternative natural goal is to maximise the intermediary's profit. This has been studied extensively starting with a paper by Myerson and Satterthwaite (Myerson and Satterthwaite 1983), which gives an analogue of Myerson's seminal result on optimal auctions, for the independent priors case. Approximately optimal mechanisms for that settings have further been studied.

(Deng et al. 2014; Niazadeh, Yuan, and Kleinberg 2014) The correlated-priors case has been investigated from a computational complexity perspective by (Gerstgrasser, Goldberg, and Koutsoupias 2016), as well as links back to auction theory (Gerstgrasser 2018). An adversarial online model of a revenue-maximising intermediation problem was studied in (Giannakopoulos, Koutsoupias, and Lazos 2017).

2 Preliminaries

In a *multi-unit bilateral trade* instance there is a buyer and seller, where the seller holds a number of units of an item. This number will be denoted by k. The buyer and seller each have a *valuation function* representing how much they value having any number of units in possession. These valuation functions are denoted by v and w, respectively. Precisely stated, a valuation function is a function $v : [k] \cup \{0\} \to \mathbb{R}_{\geq 0}$ where $v(0) = 0$. Note that we use the standard notation $[a]$, for a natural number a, to denote the set $\{1, \dots, a\}$. We denote by v the valuation function of the buyer, drawn from f, and we denote by w the valuation function of the seller, drawn from g. For $q \in [k]$, the valuation $v(q)$ or $w(q)$ of an agent (i.e., buyer or seller) expresses in the form of a number the extent to which he would like to have q units in his possession.

A mechanism $\mathbb{M}$ interacts with the buyer and the seller and decides, based on this interaction, on an *outcome*. An outcome is defined as a quadruple $(q_B, q_S, \rho_B, \rho_S)$, where q_B and q_S denote the numbers of items allocated to the buyer and the seller respectively, such that $q_B + q_S = k$. Moreover, ρ_B and ρ_S denote the payments that the mechanism charges to the buyer and seller respectively. Note that typically the payment of the seller is negative since he will get money in return for losing some items, while the payment of the buyer is positive since he will pay money in return for obtaining some items. Let $\mathcal{O}$ be the set of all outcomes. For brevity we will often refer to an outcome simply by the number of units traded q_B.

Formally, a mechanism is a function $\mathbb{M} : \Sigma_B \times \Sigma_S \to \mathcal{O}$, where Σ_B and Σ_S denote strategy sets for the buyer and seller. A *direct revelation* mechanism is a mechanism for which Σ_B and Σ_S consists of the class of valuation functions that we want to consider. That is, in such mechanisms, the buyer and seller directly report their valuation function to the mechanism, and the mechanism decides an outcome based on these reports. We want to define our mechanism in such a way that there is a dominant strategy for the buyer and seller, under the assumption that their valuation functions are in a given class $\mathcal{V}$. It is well known (see e.g. (Börgers 2015)) that then we may restrict our attention to direct revelation mechanisms in which the dominant strategy for the buyer and seller is to report the valuation functions that they hold. Such mechanisms are called *dominant strategy incentive compatible (DSIC)* for $\mathcal{V}$. In this paper, we consider for $\mathcal{V}$ two natural classes of valuation functions:

- Monotonically increasing submodular functions, i.e., valuation functions v such that for all $x, y \in [k]$ where $x < y$ it holds that $v(x) - v(x-1) \geq v(y) - v(y-1)$ and $v(x) < v(y)$. This reflects a common phenomenon observed in many economic settings involving identical goods: Possessing more of a good is never undesirable, but the increase in valuation still goes down as the held amount increases. For a monotonically increasing submodular function v and number of units $x \in [k]$, we denote by $\tilde{v}(x)$ the *marginal* valuation $v(x) - v(x-1)$. Thus, it holds that $\tilde{v}(x) \geq \tilde{v}(y)$ when $x < y$.
- Monotonically increasing functions, i.e., valuation functions v such that $v(x) < v(y)$ for all $x < y$, where $x, y \in [k]$.

Besides the DSIC requirement, there are various additional properties that we would like our mechanism to satisfy.

- Ideally, our mechanism should be *strongly budget balanced (SBB)*, which means that for any outcome $(q_B, q_S, \rho_B, \rho_S)$ that the mechanism may output it holds that $\rho_B = -\rho_S$. This requirement essentially states that all money transferred is between the buyer and the seller only.
- Additionally, we want that running the mechanism never harms the buyer and the seller. This requirement is known as *(ex-post) individual rationality (IR)*. Note that when v and w are the valuation functions of the buyer and the seller, then the initial utility of the buyer is 0 and the initial utility of the seller is $w(k)$. Thus, a mechanism $\mathbb{M}$ is individually rational if for the outcome $\mathbb{M}(v, w) = (q_B, q_S, \rho_B, \rho_S)$ it always holds that $v(q_B) - \rho_B \geq 0$ and $w(q_S) - \rho_S \geq w(k)$.
- We would like the mechanism to return an outcome for which the total utility is high. That is, we want the mechanism to maximise the sum of the buyer's and seller's utility, which is equivalent to maximizing the sum of valuations $v + w$ when strong budget balance holds.

We characterise in Section 3 the class of DSIC, SBB, IR mechanisms for both valuation classes. In the formal proof of our characterisation we describe our class in the form of direct revelation mechanisms, and we exploit the fact that we may restrict our attention to this class. However, in this version we present a more informal description of our characterization in which we decribe our class as a set of sequential posted price mechanisms, so as to provide the reader with a more intuitive understanding of the characterization.

In Section 4, we subsequently provide various approximation results on the quality of the solution output by some of these mechanisms. For these results, we assume the standard *Bayesian setting*: The mechanism has no knowledge of the buyer's and seller's precise valuation, but knows that these valuations are drawn from known probability distributions over valuation functions. Our approximation results provide mechanisms that guarantee a certain outcome quality (which is measured in terms of *social welfare*, defined in Section 4) for arbitrary distributions on the valuation functions.

Formally, in the Bayesian setting, a multi-unit bilateral trade instance is a pair (f, g, k), where $k \in \mathbb{N}$ is the total number of units that the seller initially has in his possession, and f and g are probability distributions over valuation functions of the buyer and the seller respectively. Note that we do not impose any further assumptions on these probability distributions.

3 Characterisation

In (Colini-Baldeschi et al. 2016) the authors prove that every DSIC, IR, SBB mechanism for classical bilateral trade (i.e. the case where $k = 1$) is a *fixed price mechanism*: That is, the mechanism is parametrised by a price $p \in \mathbb{R}_{\geq 0}$ such that the buyer and seller trade if and only if the buyer's valuation exceeds the price and the price exceeds the seller's valuation. Moreover, in case trade happens, the buyer pays p to the seller. In this paper we characterise the set of DSIC, IR, and WBB mechanisms for multi-unit bilateral trade, and we thereby generalise the characterisation of (Colini-Baldeschi et al. 2016).

Theorem 3.1. *Any mechanism that satisfies DSIC, IR and SBB must be a* sequential posted price mechanism with a fixed per-unit price p, potentially with bundling*, which we will refer to as a* multi-unit fixed price mechanism*. Such a mechanism iteratively proposes a quantity q of units to both the buyer and seller simultaneously, which the seller and buyer can choose to either accept or reject. If both agents accept, q additional units are reallocated from the seller to the buyer, the buyer pays pq to the seller, and the mechanism may then either proceed to the next iteration or terminate. If one of the two agents rejects, the mechanism terminates. Quantity q may vary among iterations, but must be pre-determined prior to execution of the mechanism.*

For increasing submodular valuations, any number of iterations is allowed. For general increasing valuations, the mechanism is further restricted to execute only one iterations (or equivalently, it may only offer one bundle for a fixed price).

In simple terms, our result states that for the submodular valuations case, the only thing to be done truthfully in this setting is to set a fixed per-unit price p, and ask the buyer and seller if they want to trade one or several units of the good at per-unit price p. This repeats until one agent rejects. In the general monotone case this is further restricted to a single such proposed trade. The following is a brief high-level (informally stated) argument of the proof of Theorem 3.1 for the submodular setting.

Lemma 3.2. *All prices must be fixed in advance, and cannot depend on the bid / valuation of neither the seller nor the buyer.*

Proof. This follows immediately from DSIC and SBB: By DSIC, for any outcome, the price charged to the buyer can't depend on the buyer's bid, otherwise one can construct scenarios in which the price charged by the buyer could be manipulated to the buyer's benefit by misreporting the bid. The same holds for the seller. By SBB the payment of the buyer completely determines the payment of the seller (the payment is simply negated) so neither payment can depend on either's bid. □

Theorem 3.3. *Suppose in a DSIC, SBB, IR mechanism the price for the outcome in which q units are traded is qp for a fixed per-unit price for all potential outcomes. Then the allocation chosen for a given pair of valuation functions is the one arising when asking bidders sequentially if they want to trade one unit (or a bundle of units), until one rejects.*

Proof. To see this, consider the seller's utility function $u_s(q) = q \cdot p + w(k - q)$ and the buyer's utility function $u_b(q) = v(q) - q \cdot p$, if q units would be traded at unit price p. Since both valuation functions are concave, it is easy to see that both utility functions are concave, and each has a single peak (one or more equal adjacent maxima, and no further local maxima). Furthermore they both start at 0, and once either of them becomes negative, it stays negative. Suppose we sequentially ask both bidders if they want to trade one unit for price p, until one rejects. Then the quantity traded is $\min(\text{argmax}(u_s), \text{argmax}\,(u_b))$, i.e. the first of the two peaks. If the mechanism iteratively proposes them bundles $q_1, q_2, \ldots$, then the same expression on the traded quantity would apply, but with the utility functions restricted to the domain $\{0, q_1, q_1 + q_2, \ldots\}$. If we ask them about the big all-k-item bundle, we would choose the bundle outcome iff $u(k) > u(0)$, for both, and 0 if for either of them $u(0) > u(k)$, i.e. if one (the first) of the peaks of the two utility functions restricted to $\{0, k\}$ is at 0.

Now, DSIC means that for any bid of the opposing agent, the agent cannot get anything better than what she gets by telling the truth. If the quantity traded by the mechanism would be larger than $\min(\text{argmax}(u_s), \text{argmax}\,(u_b))$, then the bidder with the lowest peak could improve her utility by claiming that all outcomes higher than her peak are wholly unacceptable (utility less than 0) to them; by IR, the mechanism would then be forced to trade the quantity at the first peak. If, on the other hand, the traded quantity would be less than the quantity of the first peak, then both players would gain by lying, in order to make the mechanism choose to trade a higher quantity (if such a quantity is at all present in the mechanism's set of tradeable quantities.) □

Theorem 3.4. *In a DSIC, SBB, IR mechanism, all potential outcomes, i.e., (quantity,price)-pairs, must have the same per-unit price.*

Proof. Suppose two outcomes have different per-unit prices. W.l.o.g. suppose for $q_1 < q_2$, $p_1/q_1 < p_2/q_2$, i.e. the per-unit price is higher in the larger allocation. Then there exists a valuation function v_{s1} for the seller in which the seller prefers outcome q_2 over q_1, but both give positive utility; and there exists another valuation function v_{s2} that gives negative utility for q_1, but the same utility for q_2. I.e. $0 < u_{s1}(q_1) < u_{s2}(q_2)$ but $u_{s2}(q_1) < 0 < u_{s2}(q_2) = u_{s1}(q_2)$. Now if for a given buyer's valuation, the chosen outcome given v_{s1} is q_1, then the seller would have an incentive to misreport v_{s2}, making outcome q_1 unavailable to the mechanism due to IR, thus making it choose q_2. Vice versa, if per-unit prices are decreasing, the same argument works for the buyer. □

Together, these three results give a full characterisation of the class of DSIC, IR, SBB mechanisms in this setting, although in our full formal proof that we provide in the full version of this paper, we need to take into account many further technical obstacles and details. There is, in particular, a *tie-breaking rule* present, that takes into account what should happen when the buyer or seller would be indifferent among multiple possible quantities, or when they would get a utility of 0 given the proposed prices and quantities.

For the case of general monotone valuations, any such mechanism must be further restricted to offering only a single outcome (other than no-trades) to the bidders.

We may alternatively view the class of mechanisms we just described as a class of direct revelation mechanisms. As direct revelation mechanisms, our precise characterisation is stated as follows.

Definition 3.1. *Let $p \in \mathbb{R}_{\geq 0}$, let $S \subseteq [k]$, and let $\tau = (\tau_B, \tau_S, \tau_\cap)$ be a vector of three tie-breaking functions specified below. The* multi-unit fixed price mechanism $\mathbb{M}_{p,S,\tau}$ *is the direct revelation mechanism that returns for a multi-unit bilateral trade instance (f, g, k) an outcome $\mathbb{M}_{p,S,\tau}(v, w) = (q_B, q_S, \rho_B, \rho_S)$ on reported valuation functions v and w, where*

- $\tau_B(v) \subseteq \arg_q \max\{v(q) - qp : q \in S \cup \{0\}\}$ *and* $\tau_B(v) \neq \varnothing$,
- $\tau_S(w) \subseteq \arg_q \max\{w(k-q) + qp : q \in S \cup \{0\}\}$ *and* $\tau_S(w) \neq \varnothing$,
- $\tau_\cap(v, w)$ *is a tie-breaking function that selects an element in $\tau_B(v) \cap \tau_S(w)$ in case this intersection is non-empty,*
- $q_B = k - q_S = \begin{cases} \min\{\max \tau_B(v), \max \tau_S(w)\} & \text{if } d_B \cap d_S = \varnothing, \\ \tau_\cap(v, w) & \text{otherwise.} \end{cases}$,
- $\rho_B = -\rho_S = q_B p$.

Informally stated, the direct-revelation version of our mechanisms offers the buyer and seller a fixed unit price p and a set of quantities S. It then asks the buyer and seller which quantity in $S \cup \{0\}$ they would like to trade at unit price p. The mechanism trades the minimum of these two demands at a unit price of p. Typically the preferred quantity is unique for both the buyer and the seller, but in case of indifferences the buyer and seller will specify a set of multiple preferred quantities. The tie-breaking functions τ_B, τ_S then determine which of these quantities are considered for trade, and $\tau_\cap$ is finally used to determine the traded quantity in case τ_B and τ_S intersect. Otherwise, the minimum of the maximum quantities of τ_B and τ_S is traded. This latter definition is the one used that allows us to give a rigorous proof of our characterisation, which we present in the full version of this paper (Gerstgrasser et al. 2018).

4 Approximation Mechanisms

In this section we study the design of DSIC, IR, SBB mechanisms that optimise the social welfare, i.e., the sum of the buyer's and seller's valuation. From Theorem 3.1, our characterization states that such a mechanism needs to be a multi-unit fixed price mechanism, so that the design challenge lies in an appropriate choice of unit-price p and quantities offered at each iteration of the mechanism.

We focus on the case of increasing submodular valuations. Obviously, every item traded can only increase the social welfare. Therefore, given that the objective is to maximise it, we repeatedly offer a single item for trade.[1] The challenge lies thus in determining the right unit price p. It is easy to see that no sensible analysis can be done if absolutely nothing is known about the valuation functions of the buyer and seller. Therefore, we assume a *Bayesian setting*, as introduced in Section 2 in order to model that the mechanism designer has statistical knowledge about the valuations of the two agents: The buyer's (and seller' valuation is assumed to be unknown to the mechanism, but is assumed to be drawn from a probability distribution f (and g) which is public knowledge. We show that we can now determine a unit price that leads to a good social welfare in expectation.

For a valuation function v of the buyer, we write $\hat{v}$ to denote the *marginal increase function* of v: $\hat{v}(q) = v(q) - v(q-1)$ for $q \in [k]$. Thus, $\hat{v}$ is a non-increasing function. Similarly, for a valuation function w of the seller, we write $\breve{w}$ to denote the *marginal decrease function* of w: $\breve{w}(q) = w(k-q+1) - w(k-q)$, for $q \in [k]$, so that $\breve{w}$ is a non-decreasing function. Thus, for all $q \in [k]$, the increase in social welfare as a result of trading q items as opposed to $q-1$ items is $\hat{v}(q) - \breve{w}(q)$. Note that therefore if v and w are increasing submodular valuation functions of the buyer and seller respectively, then the social welfare is maximised by trading the maximum number of units q such that $\hat{v}(q) > \breve{w}(q)$. We measure the quality of a mechanism on a bilateral trade instance (f, g, k) as the factor by which its expected social welfare is removed from the expected optimal social welfare $OPT(f, g, k)$ that would be attained if the buyer and seller would always trade the maximum profitable amount:

$$OPT(f, g, k) =$$

$$= \mathop{\mathbf{E}}_{v \sim f, w \sim g}\left[w(k) + \sum_{q=1}^{\max\{q' : \hat{v}(q') > \breve{w}(q')\}} (\hat{v}(q) - \breve{w}(q))\right]$$

$$= \mathop{\mathbf{E}}_{v \sim f, w \sim g}\left[\sum_{q=1}^{k} \breve{w}(q) + \sum_{q=1}^{\max\{q' : \hat{v}(q') \geq \breve{w}(q')\}} (\hat{v}(q) - \breve{w}(q))\right]$$

For $q \in [k]$ and a seller's valuation function w, we denote by $GFT(v, w, q)$ the value $\max\{0, \hat{v}(q) - \breve{w}(q)\}$ (where "GFT" is intended to stand for "Gain From Trade"). Note that $GFT(v, w, q)$ is non-increasing in q and that $OPT(f, g, k)$ can be written as

$$OPT(f, g, k) = \sum_{q=1}^{k} \mathbf{E}_{w \sim g}[\breve{w}(q) + GFT(v, w, q)].$$

Note that a social welfare as high as opt $OPT(f, g, k)$ can typically not be attained by any DSIC, IR, SBB mechanism. However, it is still a natural benchmark for measuring the performance of such a mechanism, and we will see next that there exists such a mechanism that achieves a social welfare that is guaranteed to approximate $OPT(f, g, k)$ to within a constant factor. In particular, for a mechanism $\mathbb{M}$, let $q_{\mathbb{M}}(v, w)$ be the number of items that $\mathbb{M}$ trades on reported valuation profiles (v, w), and define

$$SW(\mathbb{M}, (g, f, k)) = \mathbf{E}_{v \sim f, w \sim g}[v(q_{\mathbb{M}}(v, w)) + v(k - q_{\mathbb{M}(v,w)})]$$

[1]Also, with respect to our tie-breaking rule mentioned at the end of the last section: We simply employ the tie breaking rule that favours the highest quantity to trade, which is the dominant choice when it comes to maximising social welfare.

as the expected social welfare of mechanism $\mathbb{M}$. We say that $\mathbb{M}$ achieves an *α-approximation to the optimal social welfare*, for $\alpha > 1$, iff $OPT(g, f, k)/SW(\mathbb{M}, (g, f, k)) < \alpha$.

We show next that the multi-unit fixed price mechanism where p is set such that $\sum_{q=1}^{k} \mathbf{Pr}_{w\sim g}[\check{w}(q) \leq p] = k/2$ achieves a 2-approximation to the optimal social welfare.

Theorem 4.1. *Let (f, g, k) be a multi-unit bilateral trade instance where the supports of f and g contain only increasing submodular functions. Let $\mathbb{M}$ be the multi-unit bilateral trade mechanism where at each step one item is offered for trade at price $p = \sum_{q=1}^{k} \mathbf{Pr}_{w\sim g}[\check{w}(q) \leq p] = k/2$, until either agent reject the offer (informally: p is the price such that the seller is expected to accept to trade half of his units at price p). Mechanism $\mathbb{M}$ achieves a 2-approximation to the optimal social welfare.*

Proof. Let v be an arbitrary buyer's valuation function. We show that the mechanism achieves a 2-approximation if f is the distribution having only v in its support, and hence v is the buyer's valuation with probability 1. It suffices to prove the claim under this assumption, because the unit-price p depends on distribution g only. Hence, if $\mathbb{M}$ achieves the claimed social welfare guarantee for every fixed buyer's valuation function, then it also achieves this guarantee for every distribution on the buyer's valuation. For ease of notation, we will abbreviate $SW(\mathbb{M}, (f, g, k))$ to simply SW and we let $\ell = \max\{q : \hat{v}_k(q) \geq p\}$ be the highest quantity that the buyer would like to trade at unit-price p. In the remainder of the proof, we will omit the subscript $w \sim g$ from the expected value operator.

We first observe that SW can be written as follows, where we write $\mathbf{1}[\cdot]$ to denote the indicator function and E_q for the event that $\hat{v}(q) \geq p \geq \check{w}(q)$.

$$SW = \mathbf{E}\left[\sum_{q=1}^{k}(\check{w}(q) + \mathbf{1}[E_q]GFT(v, w, q))\right]$$
$$= \mathbf{E}\left[\sum_{q=1}^{\ell}(\check{w}(q) + \mathbf{1}[E_q]GFT(v, w, q))\right] + \mathbf{E}\left[\sum_{q=\ell+1}^{k}\check{w}(q)\right] \tag{1}$$

We will bound these last two expected values separately in terms of $OPT(f, g, k)$, and subsequently we will combine the two bounds to obtain the desired approximation factor.

We start with the quantities up to ℓ, for which first rewrite the expression as follows.

$$\mathbf{E}\left[\sum_{q=1}^{\ell}(\check{w}(q) + \mathbf{1}[E_q]GFT(v, w, q))\right]$$
$$= \sum_{q=1}^{\ell}\mathbf{E}[\check{w}(q)] + \sum_{q=1}^{\ell}\mathbf{Pr}[E_q]\mathbf{E}[GFT(v, w, q)) \mid E_q]$$
$$= \sum_{q=1}^{\ell}\mathbf{E}[\check{w}(q)] + \sum_{q=1}^{\ell}\mathbf{Pr}[E_q]\mathbf{E}[GFT(v, w, q)) \mid E_q].$$

Now, observe that $\mathbf{Pr}[E_q] = \mathbf{Pr}[p \geq \check{w}(q)]$ for quantities $q \leq \ell$. Since $\sum_{q=1}^{k}\mathbf{Pr}[p \geq \check{w}(q)] = k/2$ and $\mathbf{Pr}[p \geq \check{w}(q)]$ is decreasing in q, this implies that $\sum_{q=1}^{\ell}\mathbf{Pr}[E_q] = \sum_{q=1}^{\ell}\mathbf{Pr}[p \geq \check{w}(q)] \geq \ell/2$. Using additionally the fact that $\mathbf{E}[GFT(v, w, q)) \mid E_q]$ is also non-increasing in q, we obtain the following bound.

$$\mathbf{E}\left[\sum_{q=1}^{\ell}(\check{w}(q) + \mathbf{1}[E_q]GFT(v, w, q))\right] \tag{2}$$
$$\geq \sum_{q=1}^{\ell}\mathbf{E}[\check{w}(q)] + \frac{\sum_{q=1}^{\ell}\mathbf{Pr}[E_q]}{\ell}\sum_{q=1}^{\ell}\mathbf{E}[GFT(v, w, q)) \mid E_q]$$
$$\geq \sum_{q=1}^{\ell}\mathbf{E}[\check{w}(q)] + \frac{1}{2}\sum_{q=1}^{\ell}\mathbf{E}[GFT(v, w, q)) \mid E_q]$$
$$\geq \sum_{q=1}^{\ell}\mathbf{E}[\check{w}(q)] + \frac{1}{2}\sum_{q=1}^{\ell}\mathbf{E}[GFT(v, w, q))]$$
$$\geq \frac{1}{2}\sum_{q=1}^{\ell}\mathbf{E}[\check{w}(q) + GFT(v, w, q)] \tag{3}$$

For the quantities higher than ℓ, we first observe that non-increasingness of $\mathbf{Pr}[\check{w}(q) < p]$ in the quantity q implies that $\mathbf{Pr}[\check{w}(q) > p]$ is non-decreasing in q. Moreover, $\sum_{q=1}^{k}\mathbf{Pr}[\check{w}(q) \leq p] = k/2$ means that $\sum_{q=1}^{k}\mathbf{Pr}[\check{w}(q) > p] = \sum_{q=1}^{k}\mathbf{Pr}[\check{w}(q) \leq p]$, hence it holds that $\sum_{q=\ell+1}^{k}\mathbf{Pr}[\check{w}(q) > p] \geq \sum_{q=1}^{k}\mathbf{Pr}[\check{w}(q) \leq p]$. Therefore, we derive

$$\mathbf{E}\left[\sum_{q=\ell+1}^{k}\check{w}(q)\right] = \frac{1}{2}\sum_{q=\ell+1}^{k}\mathbf{E}[\check{w}(q)] + \frac{1}{2}\sum_{q=\ell+1}^{k}\mathbf{E}[\check{w}(q)]$$
$$\geq \frac{1}{2}\sum_{q=\ell+1}^{k}\mathbf{E}[\check{w}(q)] + \frac{1}{2}\sum_{q=\ell+1}^{k}\mathbf{E}[\check{w}(q) \mid \check{w}(q) > p]\mathbf{Pr}[\check{w}(q) > p]$$
$$\geq \frac{1}{2}\sum_{q=\ell+1}^{k}\mathbf{E}[\check{w}(q)] + \frac{1}{2}\sum_{q=\ell+1}^{k}\hat{v}(q)\mathbf{Pr}[\check{w}(q) > p]$$
$$\geq \frac{1}{2}\sum_{q=\ell+1}^{k}\mathbf{E}[\check{w}(q)] + \frac{1}{2}\sum_{q=\ell+1}^{k}\mathbf{E}[GFT(v, w, q)]$$
$$\geq \frac{1}{2}\sum_{q=\ell+1}^{k}\mathbf{E}[\check{w}(q) + GFT(v, w, q)], \tag{4}$$

where the second inequality holds because $\check{w}(q)$ conditioned on $\check{w}(q) > p$ is always higher than $\hat{v}(q)$ which does not exceed p. Moreover, the third inequality follows because $\mathbf{E}[GFT(v, w, q)] = \mathbf{E}[(\hat{v}(q) - \check{w}(q))\mathbf{1}(\hat{v}(q) > \check{w}(q))] \leq \mathbf{E}[\hat{v}(q)\mathbf{1}(\hat{v}(q) > \check{w}(q))] \leq \mathbf{E}[\hat{v}(q)\mathbf{1}(p > \check{w}(q))] = \hat{v}(q)\mathbf{Pr}[p > \check{w}(q)]$.

We now use (3) and (4) to bound (1) and obtain the desired inequality

$$SW \geq \frac{1}{2}\sum_{q=1}^{k}\mathbf{E}[\check{w}(q) + GFT(v, w, q)] = \frac{OPT(f, g, k)}{2},$$

which proves the claim. □

The above 2-approximation mechanism is deterministic. We show next that we can do better if we allow randomisation: Consider the *Generalized Random Quantile Mechanism*, or $\mathbb{M}_G$, which draws a number x in the interval $[1/e, 1]$

where the CDF is $\ln(ex)$ for $x \in [1/e, 1]$. The mechanism then sets a unit price $p(x)$ such that $\mathbf{E}_w[\max\{q : w(q) \geq qp(x)\}] = \sum_{q=1}^{k} \mathbf{Pr}_w[\breve{w}(q) \leq p(x)] = xk$, repeatedly offering single item trades as before. In words, the price is set such that the expected number of units that the seller is willing to sell, is an x fraction of the total supply, where x is randomly drawn according to the probability distribution just defined. This randomised mechanism satisfies DSIC, IR, and SBB, because it is simply a distribution over multi-unit fixed price mechanisms. Note that this mechanism is also a generalisation of a previously proposed mechanism: In (Blumrosen and Dobzinski 2016), the authors define the special case of this mechanism for a single item, and call it the *Random Quantile Mechanism*. They show that it achieves a $e/(e-1)$-approximation to the social welfare, and we will prove next that this generalisation preserves the approximation factor, although the proof we provide for it is substantially more complicated and requires various additional technical insights.

Theorem 4.2. *Let (f, g, k) be a multi-unit bilateral trade instance where the supports of f and g contain only increasing submodular functions. The Generalised Random Quantile Mechanism $\mathbb{M}_G$ achieves a $e/(e-1)$-approximation to the optimal social welfare.*

Proof sketch. As in the proof of Theorem 4.1, we fix a valuation function v for the buyer. The proof works by first rewriting $OPT(f, g, k)$ as follows:

$$OPT(f,g,k) = \sum_{q=1}^{k} \hat{v}(q) + \sum_{q=1}^{k} (\mathbf{E}_w[\breve{w}(q) \mid \breve{w}(q) \geq \hat{v}(q)] - \hat{v}(q)) \mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(b)]. \quad (5)$$

Then, the proof proceeds by deriving a lower bound of $(1 - 1/e)$ times the expression (5) on SW, which implies our claim. It can be derived that SW can be bounded and rewritten into a sum of three separate summations over the items.

$$SW \geq \sum_{q=1}^{k} \hat{v}(q) \mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] + \sum_{q=1}^{k} \hat{v}(q) \mathbf{Pr}_{w,x}[p(x) \in [\breve{w}(q), \hat{v}(q)] \mid \breve{w}(q) < \hat{v}(q)] \cdot \mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)] \quad (6)$$
$$+ \sum_{q=1}^{k} (\mathbf{E}_w[\breve{w}(q) \mid \breve{w}(q) \geq \hat{v}(q)] - \hat{v}(q)) \mathbf{Pr}[\breve{w}(q) \geq \hat{v}(q)].$$

Next, we bound the first part (6) of the last expression, i.e., excluding the last summation.

$$
\begin{aligned}
(6) &\leq \sum_{q=1}^{k} \hat{v}(q) \mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] + \sum_{q=1}^{k} \hat{v}(q) \mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)] \cdot \frac{\int_{1/e}^{z:p(z)=\hat{v}(q)} \mathbf{Pr}_w[\breve{w}(q) \leq p(x)] \frac{1}{x} dx}{\mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)]} \\
&= \sum_{q=1}^{k} \hat{v}(q) \mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] + \int_{1/e}^{z:p(z)=\hat{v}(q)} \left(\sum_{q=1}^{k} \hat{v}(q) \mathbf{Pr}_w[\breve{w}(q) \leq p(x)] \right) \frac{1}{x} dx \\
&\geq \sum_{q=1}^{k} \hat{v}(q) \mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] + \int_{1/e}^{z:p(z)=\hat{v}(q)} \sum_{q=1}^{k} \hat{v}(q) \frac{kx}{k} \frac{1}{x} dx \\
&= \sum_{q=1}^{k} \hat{v}(q) \mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] + \sum_{q=1}^{k} \hat{v}(q) \int_{1/e}^{z:p(z)=\hat{v}(q)} 1 dx \\
&= \sum_{q=1}^{k} \hat{v}(q) \mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] + \sum_{q=1}^{k} \hat{v}(q) (\mathbf{Pr}[\breve{w}(q) < \hat{v}(q)] - \frac{1}{e}) \\
&= (1 - 1/e) \sum_{q=1}^{k} \hat{v}(q), \qquad (7)
\end{aligned}
$$

where for the inequality we used that both $\hat{v}(q)$ and $\mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)]$ are non-increasing in q, so that replacing all the probabilities by the average probability xk/k yields a lower value. Substituting (6) by (7) and using the expression (5) for OPT then yields the desired bound.

$$
\begin{aligned}
SW &\geq (1 - 1/e) \Big(\sum_{q=1}^{k} \hat{v}(q) + \sum_{q=1}^{k} (\mathbf{E}_w[\breve{w}(q) \mid \breve{w}(q) \geq \hat{v}(q)] - \hat{v}(q)) \mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] \Big) \\
&= (1 - 1/e) OPT(f, g, k).
\end{aligned}
$$

□

For a rigorous proof of this result, we refer the reader to the full version of this paper (Gerstgrasser et al. 2018).

Currently we have no non-trivial lower bound on the best approximation factor achievable by a DSIC, IR, SBB mechanism, and we believe that the approximation factor of $e/(e-1)$ achieved by our second mechanism is not the best possible. For our first mechanism, it is rather easy to see that the analysis of the approximation factor of our first mechanism is tight, and that it is a direct extension of the median mechanism of (McAfee 2007), for which it was already shown in (Blumrosen and Dobzinski 2016) that it does not achieve an approximation factor better than 2: The authors show that 2 is the best approximation factor possible for any deterministic mechanism for which the choice of p does not depend on the buyer's distribution.

For the more general class of increasing valuation functions, an approximation factor of $(2e-1)/(e-1) \approx 2.582$ to the optimal social welfare is achieved by a mechanism of (Blumrosen and Dobzinski 2016): They use a $e/(e-1)$-approximation mechanism for the single-item setting, which yields a $(2e-1)/(e-1)$ approximation mechanism for the multi-unit setting through a conversion theorem which they prove. We note that their conversion theorem is more precisely presented for the setting with a buyer and a seller who holds one *divisible* item. However, their proof straightforwardly carries over to the multi-unit setting. It would be an interesting open challenge to improve this currently best-known bound of $(2e-1)/(e-1)$ for general increasing valuations.

5 Discussion

Our results give a full characterisation of truthful mechanisms for the multi-unit bilateral trade setting. This is of importance not only for theoretical considerations, but also due to its practical consequences: We have shown that the class of truthful mechanisms in this setting consists only of very simple constant unit-price, sequential posted price mechanisms. These are not only obviously truthful, but also very easy to implement. They require little computation on the participants' side, and the communication complexity of such a protocol is minimal.

Many interesting open questions remain in this area. In the simple setting we consider, we do not know matching upper and lower bounds on the approximation ratio. For the multi-unit setting studied here, the next step would be to generalize first to markets with multiple buyers and sellers and then to indistinguishable agents, each entering the market with an endowment of items. In an orthogonal direction, it is interesting to consider the case with a single pair of a seller and a buyer, but multiple item types that are substitutes of each other, with more complex valuation functions as described in e.g. (Kelso Jr and Crawford 1982).

References

Babaioff, M.; Cai, Y.; Gonczarowski, Y. A.; and Zhao, M. 2018. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *Procs. of the 19th ACM Conference on Economics and Computation*, 373–373.

Balseiro, S.; Mirrokni, V.; Leme, R. P.; and Zuo, S. 2018. Dynamic double auctions: Towards first best. *SSRN* 3213460:38pp.

Blumrosen, L., and Dobzinski, S. 2016. (almost) efficient mechanisms for bilateral trading. *ArXiv preprint, arXiv:1604.04876*.

Blumrosen, L., and Mizrahi, Y. 2016. Approximating gains-from-trade in bilateral trading. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*, 400–413. Springer.

Börgers, T. 2015. *An Introduction to the Theory of Mechanism Design*. Oxford University Press.

Brustle, J.; Cai, Y.; Wu, F.; and Zhao, M. 2017. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, 589–590. ACM.

Colini-Baldeschi, R.; de Keijzer, B.; Leonardi, S.; and Turchetta, S. 2016. Approximately efficient double auctions with strong budget balance. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1424–1443. SIAM.

Colini-Baldeschi, R.; Goldberg, P.; de Keijzer, B.; Leonardi, S.; and Turchetta, S. 2017a. Fixed price approximability of the optimal gain from trade. In R. Devanur, N., and Lu, P., eds., *Procs. of the 13th conference on Web and Internet Economics*, 146–160. Springer International Publishing.

Colini-Baldeschi, R.; Goldberg, P. W.; de Keijzer, B.; Leonardi, S.; Roughgarden, T.; and Turchetta, S. 2017b. Approximately efficient two-sided combinatorial auctions. In *Proceedings of the 18th ACM Conference on Economics and Computation*, EC '17, 591–608. ACM.

Deng, X.; Goldberg, P.; Tang, B.; and Zhang, J. 2014. Revenue maximization in a bayesian double auction market. *Theoretical Computer Science* 539:1–12.

Feldman, M., and Gonen, R. 2016. Online truthful mechanisms for multi-sided markets. *arXiv preprint, arXiv:1604.04859*.

Feldman, M., and Gonen, R. 2018. Removal and threshold pricing: Truthful two-sided markets with multi-dimensional participants. In Deng, X., ed., *Proceedings of the Symposium on Algorithmic Game Theory (SAGT)*, 163–175. Springer.

Gerstgrasser, M.; Goldberg, P. W.; de Keijzer, B.; Lazos, P.; and Skopalik, A. 2018. Multi-unit bilateral trade. *arXiv preprint, arXiv:1811.05130*.

Gerstgrasser, M.; Goldberg, P. W.; and Koutsoupias, E. 2016. Revenue maximization for market intermediation with correlated priors. In *International Symposium on Algorithmic Game Theory*, 273–285. Springer.

Gerstgrasser, M. 2018. On the complexity of optimal correlated auctions and reverse auctions. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 605–613. International Foundation for Autonomous Agents and Multiagent Systems.

Giannakopoulos, Y.; Koutsoupias, E.; and Lazos, P. 2017. Online market intermediation. *arXiv preprint, arXiv:1703.09279*.

Kelso Jr, A. S., and Crawford, V. P. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society* 1483–1504.

McAfee, R. P. 2007. The gains from trade under fixed price mechanisms. Technical report, California Institute of Technology.

Myerson, R. B., and Satterthwaite, M. A. 1983. Efficient mechanisms for bilateral trading. *Journal of economic theory* 29(2):265–281.

Niazadeh, R.; Yuan, Y.; and Kleinberg, R. 2014. Simple and near-optimal mechanisms for market intermediation. In *International Conference on Web and Internet Economics*, 386–399. Springer.

Segal-Halevi, E.; Hassidim, A.; and Aumann, Y. 2018a. Double auctions in markets for multiple kinds of goods. In *Proceedings of the 27th IJCAI Conference*, 489–497.

Segal-Halevi, E.; Hassidim, A.; and Aumann, Y. 2018b. Muda: A truthful multi-unit double-auction mechanism. In *Proceedings of the 32nd AAAI Conference*, 1193–1201.