Combinatorial Reallocation Mechanisms

Liad Blumrosen* Shahar Dobzinski

August 20, 2018

Abstract

We consider reallocation problems in settings where the initial endowment of each agent consists of a subset of the resources. The private information of the players is their value for every possible subset of the resources. The goal is to redistribute resources among agents to maximize efficiency. Monetary transfers are allowed, but participation is voluntary.

We develop incentive-compatible, individually-rational and budget-balanced mechanisms for two settings in which agents have complex multi-parameter valuations, both settings include double auctions as a special case. The first setting is combinatorial exchanges, where we provide a mechanism that achieves a logarithmic approximation to the optimal efficiency when valuations are subadditive. The second setting is Arrow-Debreu markets for a single divisible good, where we present a constant approximation mechanism. The first result is given for a Bayesian setting, where the latter result is for prior-free environments.

*A preliminary version of the results in this paper was presented in ACM EC 2014 under the title "Reallocation Mechanisms" and appeared in a 1-page abstract.

1 Introduction

A fundamental problem in economics is the allocation of scarce resources. Initially, resources may be inefficiently distributed among agents. However, as agents value resources differently, trade might improve their well being.

When each agent seeks to maximize his own utility, classic economic theory generally predicts the existence of an "invisible hand": agents will trade among themselves to maximize their own utility, and will eventually arrive at an efficient resource allocation. This paradigm fails, however, in the presence of asymmetric information. Private information may lead to market failures, where trade does not take place even when it is desirable for all. One influential example is Akerlof's market for used cars [1], where the unraveling of markets leads to no trade at all. In this paper we aim to design mechanisms that will foster trade in such markets, even if resources are distributed among multiple agents with different interests and partial information.

We focus on exchange economies, where each agent is initially endowed with some resources and thus agents simultaneously play the role of buyers and sellers. Exchange economies are related to many real life environments; Individuals hold assets, like real estate, cars and stocks, for which other individuals have their own preferences. Firms hold other types of assets (e.g., employees, land, machines, intellectual property) that may possibly be better assigned if more information becomes available. Numerous examples from the realms of industrial organization and finance fall into this model, like the dissolving of partnerships, breaking monopolies, merges and acquisitions and other anti-trust related acts. Of particular interest are structured markets where trade can be coordinated by centralized mechanisms. One recent example for a centralized large-scale reallocation mechanism is the FCC's attempt to reallocate frequencies currently held by TV broadcasters to wireless phone companies (see [23]). A major challenge in these FCC two-sided auctions is to provide incentives for the TV broadcasters to relinquish their licenses (see also [3]).

A simpler version of exchange markets are two-sided markets (also known as *double auctions*, see, e.g., [28, 29, 12, 20, 19]). In such settings, agents are classified in advance as buyers and sellers, buyers cannot sell anything and sellers cannot purchase additional items. Two-sided markets received much attention recently with the emergence of the *sharing economy*, where many large scale two-sided platforms became highly popular (examples include Airbnb for short-term lodging, Uber for transportation and Prosper for personal loans). Our results thus also hold for two sided markets, which is a special case of our more general exchange framework.

The main goal of this paper is to design markets that enable efficient reallocation of resources. Technically, this translates to three requirements. The first one is *individually rationality*: the participation of the agents is voluntary and at any point they may leave the market and consume their initial endowments. Thus, the outcome of any individually rational mechanism is a Pareto improvement in the economy, where agents are expected to be (weakly) better off in the new allocations. The second requirement is *budget balance*: the mechanism is not allowed to subsidize the agents in order to improve the outcome. We distinguish between *weakly* and *strongly* budget balanced mechanisms: in the latter the mechanism is additionally not allowed to burn money[1].

The third requirement is *truthfulness*. We discuss both Bayesian and prior-free models, but all our mechanisms admit ex-post dominant-strategy equilibria. Even when distributional assumptions are made, we make minimal use of this knowledge, namely we only require access to statistical properties like the endowment's median value[2].

[1]The budget balance requirement is common in the cost-sharing literature (e.g., [24] and [27]) but there the idea is to charge the participants an amount that suffices to cover the cost of providing the service.

[2]In fact, using noisy estimations of the medians decreases the performance of our mechanism in a rate proportional to the noise. Hence, even if we only have a black box access to the distributions, we can use the black box to estimate

Our reallocation problems are essentially combinatorial auctions where items are initially held by the players (and not by the auctioneer as usual), hence generalizing models of double auctions. This adds another layer of complexity; For example, while VCG (Vickrey-Clarke-Groves, see [26]) mechanisms can always be used to maximize welfare in combinatorial auctions, in the presence of endowments no truthful mechanism can allocate efficiently and remain budget balanced, as we will shortly see.

While several recent papers considered two-sided markets with single-parameter preferences ([5, 6, 30, 10, 7]), we consider agents that have complex *multi-parameter* valuations. Designing mechanisms for multi-parameter settings is one of the main challenges of Algorithmic Mechanism Design. The additional properties we require (e.g., budget balance) make the design of mechanisms in our case is even more demanding.

We consider two main settings in this paper. In the first, agents have combinatorial valuations, where valuations are drawn from distributions that are known to the market planner. We then explore prior-free environments where agents want to exchange quantities of one, fully divisible good; Their valuations for the good can be described by any concave real function.

Combinatorial Exchanges

In a *combinatorial exchange* there are n players and a set of M indivisible, heterogeneous items ($|M| = m$). Each player i has a valuation function $v_i : 2^M \to \mathbb{R}^+$ that specifies his value for each possible subset of the items. We assume that each i is non-decreasing and that $v_i(\emptyset) = 0$. Crucially, we assume that the items are initially distributed among the agents. That is, each agent i holds a (possibly empty) subset E_i of the items.

Recall that we are looking for truthful mechanisms that reallocate the items in a way that improves the welfare. Our mechanism should be individually rational, i.e., the utility of each bidder (from the new allocation + payment) should not be smaller than his value for his initial endowment. Of course, we avoid subsidies by requiring that the mechanism is weakly budget balanced.

Let H_n denote the n'th harmonic number, and let $t = \max_i |E_i|$ (i.e., the maximal number of items held by a single player). We are able to show that:

Theorem: There exists a truthful, individually rational, weakly budget balanced, randomized mechanism that provides an $8 \cdot H_t$-approximation to the optimal welfare if all valuations are subadditive[3]. The only distributional knowledge that the mechanism uses is the median value of the distribution of the endowment of each player.

In particular, if each bidder is initially endowed with at most one item we get an 8-approximation.

To gain some intuition about the mechanism, let us consider a simpler setting in which one agent i initially holds all items (that is, $E_i = M$). Let MED_i denote the median of the value of the distribution of $v_i(E_i)$. Consider the following (incorrect) mechanism: agent i reports whether he agrees to sell all items for a price of MED_i. If so, we will use VCG to find an optimal allocation of the items to all agents but agent i, and use the revenue generated from VCG to pay an amount of MED_i to agent i. If not, agent i keeps the items. Notice that the approximation ratio of this (incorrect) procedure is constant: if most of the expected optimal welfare is contributed by agent i then by doing nothing we already get a 2 approximation, and the outcome of any valid mechanism is a Pareto improvement. On the other hand, if most of the expected optimal welfare is contributed

the medians within an arbitrary precision, and preserve very similar performance guarantees.

[3]A valuation v is subadditive if for ever two bundles S and T we have that $v(S) + v(T) \geq v(S \cup T)$.

by the rest of the agents, then we get a 4 approximation: with probability $\frac{1}{2}$ agent i sells his endowment, and in that case we allocate the items optimally among all agents but agent i.

Of course, the procedure above fails because we cannot guarantee that the revenue of the VCG mechanism will be at least MED_i. To handle this, we develop a "revenue extracting" procedure which is the combinatorial-auctions analogue of a second-price auction with a reserve price. In a second-price auction the auctioneer can put a reserve price r to guarantee revenue of at least r when the highest value is at least r. We show that in a combinatorial auction with n players there exists a (deterministic, prior-free) mechanism that guarantees a revenue of at least r if the optimal welfare is at least $H_n \cdot r$.[4]

Our mechanism (for the special case) now works as follows: allow agent i to sell all his items at price MED_i. If agent i agrees, we use the revenue extraction procedure with a "global reserve price" $r = MED_i$ to distribute the items among all agents but agent i. We then use the revenue to pay MED_i to agent i. The mechanism for the general case can be found in Section 3.

Prior-Free Mechanisms: Arrow-Debreu Markets

Our second main technical construction considers the classic exchange model of Arrow and Debreu [2]. We have a single divisible good and the valuations can be any function with decreasing marginals (i.e., concave valuations). An easy adaption of the mechanism for combinatorial exchanges guarantees a constant approximation, but the challenge now is to get rid of the distributional assumptions and develop *prior free* mechanisms with a constant approximation ratio in the worst case.

From a technical perspective this is a multi-parameter environment for which our machinery for developing truthful mechanisms (especially prior free ones) is limited. Yet, to our surprise we were able to come up with a prior-free constant-approximation mechanism:

Theorem: There exists a truthful, prior-free, individually rational, weakly budget balanced, randomized mechanism that provides a constant approximation to the optimal social welfare, as long as no player is initially endowed with more than $\frac{1}{8}$ of the good.

Notice the necessity of the last condition: in markets when, say, one player initially holds all the good, in the spirit of [25, 21] no prior-free mechanism with a constant approximation ratio exists.

A key idea in the mechanism is to replace the revenue extraction procedure that was used in the mechanism for combinatorial exchanges with a more subtle one. The new procedure allows us to take advantage of the specifics of the setting and get a constant approximation ratio. Consider the special case where one agent i initially holds the whole good. Now, since we assume no distributional knowledge, we do not know the median value of the endowment of i, but let us assume for now that we know instead the "mid-supply" price: the price (per fraction) p for which i prefers to sell exactly half of his endowment. The crux is that since the valuation of i exhibits decreasing marginals, agent i will agree to sell *any* amount smaller than half of his endowment at the price. Thus, if we knew that mid-supply price we could just run VCG with the rest of the agents as well as an additional dummy buyer that has a value (per fraction) of p for any amount below half of the endowment of i. Observe that if some agent i' was allocated fraction x of the good, the VCG payment formula implies that his payment his at least $x \cdot p$ (since otherwise the dummy buyer can get an additional amount x of the good). We can now take an amount of x from agent i, assign it to i' and pay agent i $x \cdot p$. The formal mechanism for the general case and its analysis are presented in Section 4.

[4]If the optimal welfare is smaller than $H_n \cdot r$, then the mechanism is not required to allocate the items, but if it does so the revenue is guaranteed to be at least r.

Conclusions and Future Directions

In this paper we devise welfare-maximizing reallocation mechanisms. We do not know whether the approximation factors achieved by our mechanisms are optimal. In particular, proving impossibilities on the power of truthful and budget balanced mechanisms for reallocation problems is an interesting open question.

Our focus in this paper was not computational complexity, but it turns out that our mechanism for Arrow-Debreu markets does run in polynomial time. However, our mechanism for combinatorial exchanges does not (see [22, 4] for computational issues in combinatorial exchanges). Developing a polynomial time mechanism for the latter setting seems hard as in particular it implies a solution to the notorious problem of developing truthful polynomial time algorithm for combinatorial auctions with subadditive (and submodular) bidders (see, e.g., [13, 18, 17, 14]).

Several follow-ups to an earlier version of this paper studied approximation mechanisms for two-sided markets. Colini-Baldeschi et al. [8] gave a constant factor approximation to several two-sided market Bayesian settings with a single-item supply sellers and unit demand buyers. The mechanisms in [8] are strongly budget balanced, while the mechanisms in our paper are weakly budget balanced (i.e., may leave leftover money). Combinatorial two-sided markets were discussed in [9], where an improved notion for strong budget balance was introduced and used for agents with restricted subadditive valuations (that is, XOS and additive). Approximately efficient mechanisms in prior-free two-sided market with identical goods were discussed in [30].

In a different direction, our technique for "combinatorial auction for a global reserve price" was extended to *combinatorial cost sharing* [16]: a cost sharing model where there are multiple non-identical resources and players have combinatorial preferences over them.

Organization

Section 2 defines a general framework which captures all settings we discuss. We study combinatorial exchanges in Section 3 and Arrow-Debreu markets in Section 4.

2 The Framework

Consider a set of resources $M = \{1, ..., m\}$ and a set of n agents. Let $\mathcal{E}_i \subseteq \{0,1\}^m$ be the set of allowed endowments for agent i. Let $\mathcal{A} \subseteq \{0,1\}^{n\times m}$ be the set of allowed allocations of resources to the agents, where $\mathcal{A}_i \subseteq \{0,1\}^m$ stands for the set of possible allocations to player i.

The valuation of player i is a function $v_i : \mathcal{A}_i \cup \mathcal{E}_i \to \mathbb{R}_+$. Let V_i be the set of all possible valuations of player i, and $V = V_1 \times ... \times V_n$. We sometimes assume a Bayesian model, where v_i is drawn from V_i according to a distribution F_i, independently from the valuations of the other agents. The valuations are private information and the endowments are known to the designer.

A (direct revelation) reallocation mechanism consists of an allocation function $M : V \to \Delta(\mathcal{A})$ and a payment function $p : V \to \mathbb{R}^n$. As agents in our model can be sellers and buyers simultaneously, we do not assume that payments are positive; Negative payments mean transfers from the mechanism to the agents.

All of our mechanisms are dominant-strategy truthful. That is, reporting the true valuations v_i is a dominant strategy for every agent i. Truthful behavior is ex-post (rather than dominant-strategy in expectation) which allows us, e.g., to ignore distributional beliefs of the agents and whether they are risk-neutral or not. We require the following:

- **Ex Post Individual Rationality.** For every allocation and payment (A_i, p_i) eventually allocated to agent i with initial endowment $E_i \in \mathcal{E}_i$ (after the realization of the valuations and the randomness of the mechanism), we have that $v_i(A_i) - p_i \geq v_i(E_i)$.

- **Ex Post Budget Balance.** For every $\mathbf{v} \in V$ of the preferences we have $\sum_{i=1}^{n} p_i(\mathbf{v}) = 0$. If, instead, we only have that $\sum_{i=1}^{n} p_i(\mathbf{v}) \geq 0$, the mechanism is *weakly budget balanced*.[5]

- **Approximate Efficiency**. We would like to approximate the optimal expected efficiency with non-strategic agents, $OPT = \max_{A \in \mathcal{A}} E_{\mathbf{v} \in V}\left[\sum_{i=1}^{n} v_i(A_i)\right]$. A mechanism achieves an α approximation to the optimal welfare if $E[\sum_{i=1}^{n} v_i(M(\mathbf{v}))] \geq \frac{OPT}{\alpha}$ (expectation is over the random coins of the mechanism, if any, and over the valuations $\mathbf{v} \in V$).

3 Combinatorial Exchanges

In this section we consider *combinatorial exchanges*: there are n agents, each agent i initially holds a subset E_i of the items. Items are heterogeneous and indivisible. Every agent i has a subadditive valuation v_i, that is, for every two bundles S, T we have that $v_i(S \cup T) \leq v_i(S) + v_i(T)$. Each v_i is independently drawn from a distribution F_i. However, our mechanism will only require that the mechanism knows, for each agent i, the median value MED_i for the bundle E_i she initially owns.

Let H_n be the n'th harmonic number ($H_n = \sum_{i=1}^{n} \frac{1}{i}$) and $t = \max_i |E_i|$. We present a mechanism that achieves an $8H_t$ approximation in this multi-parameter domain. Importantly, if each player is initially endowed with at most one item we get an 8-approximation.

In the first step of the mechanism, the bidders are randomly partitioned into two sets, "buyers" and "sellers", and each "seller" i is offered to sell his endowment bundle at a price MED_i. Then we would like to take all the items that were sold and optimally allocate them among the "buyers" using VCG. The main obstacle is that VCG is not budget balanced. To overcome this we present a procedure that guarantees (approximate) welfare maximization while guaranteeing some minimal amount of revenue. Subsection 3.1 describes this procedure and the mechanism itself is in Subsection 3.2.

3.1 Detour: Combinatorial Auctions with Global Reserve

In this subsection, we depart for a while from the general exchange model and study a standard one-sided combinatorial auction model. Our mechanism for combinatorial exchanges will use the insights obtained in this subsection.

Consider a standard combinatorial auction with a set M of m heterogeneous items that has to be allocated to n bidders. Each bidder i has a valuation $v_i : 2^M \to \mathbb{R}$. As usual we assume that each v_i is normalized ($v_i(\emptyset) = 0$) and non-decreasing. While the standard goal in the literature is to maximize welfare, assuming the auctioneer has no production cost for the items, in our case the auctioneer has a non-negative cost r of producing all items[6]. We are interested in truthful and individually rational mechanisms.

We first observe that there are some simple cases where the auctioneer cannot cover the production cost. Denote by OPT the value of a welfare-maximizing allocation of the items in M to the bidders. Therefore, the interesting case is when $OPT > r$. Namely, given $\alpha \geq 1$, whenever $OPT \geq \alpha \cdot r$ the mechanism must allocate some items to the bidders and raise a revenue of at least

[5]Note that this definition holds for every realization of v (and not only in expectation, which is usually a key for achieving budget balance in Bayesian domains, e.g., in [11]).

[6]The items are produced only if a sale is made. Here we assume for simplicity that the cost of producing the first item is r and the cost of producing any additional item is 0. This corresponds to the case that the production cost of items is governed by the start-up cost. A more realistic setup assumes a production cost for each item, or more generally for bundles of items. Indeed, the mechanism of Subsection 3.2 essentially provides a solution for this case.

r. Else, when $OPT < \alpha \cdot r$ the mechanism is not required to sell the items (but, again, if it does sell the revenue must be at least r).

The challenge is, of course, to develop such a mechanism with α that is as small as possible. We do so for $\alpha = H_n$. We achieve this by relying on the well known observation that VCG generalizes to maximization of an "adjusted welfare" according to an affine function (see, e.g., [26] for further details). Specifically, instead of maximizing social welfare, $\Sigma_i v_i(A_i)$, we run a VCG mechanism that selects an allocation with the highest "adjusted welfare": we "adjust" the welfare of an allocation $A = (A_1, \ldots, A_n)$ to be $\Sigma_i v_i(A_i) - H_{n_A} \cdot r$, where n_A is the number of non-empty bundles in the allocation A.

If we use VCG payments, we obtain a truthful mechanism. In addition, the mechanism allocates some items whenever $OPT \geq \alpha \cdot r$. To see this, observe that the mechanism allocates some items only if there is an allocation with a positive adjusted welfare (since the adjusted welfare of the empty allocation is 0). Thus, all that is left is to prove that when the mechanism allocates some items then the revenue is at least r. Suppose that the mechanism outputs the allocation $A = (A_1, \ldots, A_n)$. Consider some bidder i with $A_i \neq \emptyset$. In the VCG mechanism bidder i pays his "damage to society". Observe that the damage to society of bidder i is at least $\frac{r}{n_A}$: consider the allocation A' that allocates each bidder $i' \neq i$ the same set of items and allocates nothing to bidder i. If we ignored the preferences of bidder i, we could have chosen the allocation A' and increase the adjusted welfare by

$$(\Sigma_{i' \neq i} v_{i'}(A_{i'}) - H_{n_A - 1} \cdot r) - (\Sigma_{i' \neq i} v_{i'}(A_{i'}) - H_{n_A} \cdot r) = H_{n_A} \cdot r - H_{n_A - 1} \cdot r = \frac{r}{n_A}$$

I.e., the VCG payment of each bidder i with $A_i \neq \emptyset$ is at least $\frac{r}{n_A}$. Since by definition there are n_A such bidders, the total revenue is at least r, as required.

We conclude that we can design a dominant-strategy truthful mechanism in which a seller can extract revenue of at least the global reserve r from n buyers whenever the optimal welfare is greater than r by a logarithmic factor, that is, whenever $OPT > H_n \cdot r$.

We note that follow-up work has shown connections between our mechanism for combinatorial auctions with global reserve and cost sharing. In particular, the mechanism that we provide here gives almost immediately an H_n approximation to the social welfare for the excludable public good problem. Furthermore, the paper [16] extends our mechanism for more general cost sharing settings. We refer the interested reader to [16] for more details.

3.2 The Combinatorial Median Mechanism

We can now finally present the *Combinatorial Median Mechanism* for combinatorial exchanges. In the first step, we randomly assign agents to roles of sellers and buyers; Since the agents have sub-additive preferences, we will show that we lose a constant factor in the approximation by doing that. We then choose sellers with low values (those who reject prices equal to their median values) and sell their items via a VCG mechanism. We use a methodology that builds on combinatorial auctions with global reserves to guarantee budget balance. The mechanism runs as follows:

1. Each player is assigned to either group S or group B uniformly at random.

2. Each player $i \in S$ will be offered a price equal to MED_i (i.e., the median value of her endowment). Let $\hat{S}$ denote the set of players in S that accepted the price. The total set of items of players in $\hat{S}$ endowments is denoted by $E_{\hat{S}} = \cup_{i \in \hat{S}} E_i$.

3. Given an allocation A of items in $E_{\hat{S}}$ to players in B, denote for each $i \in \hat{S}$ by t_i the number of buyers that hold in A at least one item from E_i, i.e., $t_i = |\{j|A_j \cap E_i \neq \emptyset\}|$. Let $c_A = \Sigma_{i\in\hat{S}} H_{t_i} \cdot MED_i$.

4. Run a VCG auction for the items $E_{\hat{S}}$ among the bidders B where we penalize the welfare of an allocation A by c_A, taking into account the endowments of bidders in B. I.e., we find the allocation A that maximizes: $\Sigma_{i\in B} v_i(A_i + E_i) - c_A$.

5. Consider seller $i \in \hat{S}$. If at least one item from his endowment E_i is sold in the VCG auction, then i is paid MED_i and loses all his endowment. Else, seller i keeps his endowment and is not paid anything. Each buyer is allocated the items he won in the VCG auction (in addition to his endowment) and pays his VCG payment.

Generally speaking, the role of the c_A's is to guarantee a lower bound to the revenue, just as in the combinatorial auction with global reserve. However, the analysis is slightly more complicated. Recall that t is the maximal endowment size, i.e., $t = \max_i |E_i|$.

Theorem 3.1. *The Combinatorial Median mechanism provides an approximation ratio of* $8 \cdot H_t$ *to the optimal social welfare. It is truthful, ex-post individually rational and weakly budget balanced.*

Proof. In the analysis we use the following simple observation:

Claim 3.2. *Let v be a subadditive valuation and S a set. Let $S_1, \ldots, S_t$ be a partitioning of S such that $\cup_r S_r = S$ and $S_r \cap S_{r'} = \emptyset$, for every $r \neq r'$. Let T be the bundle that is composed by adding each S_r independently at random, with probability at least p. Then, $E[v(T)] \geq p \cdot v(S)$.*

Proof. (of Claim 3.2) We will use the following folklore observation:

Observation 3.3. *Let v be a subadditive valuation and S a set. Suppose that each item in S is added to a bundle T independently with probability at least p. Then, $E[v(T)] \geq p \cdot v(S)$.*

Given v, let M' be a set of t items, where each item j corresponds to S_j. We will define a valuation function v' on M' in the following way: for each $U \subseteq M'$, $v'(U) = v(\cup_{j\in U} S_j)$. Observe that selecting items at random from M' and measuring their value according to v' is exactly like selecting the bundles $S_1, \ldots, S_t$ at random and measuring their value according to v. Thus, since v' is subadditive whenever v is subadditive we have that the statement of the claim follows, simply by applying the observation to v. □

Fix some welfare maximizing solution $(O_1, \ldots, O_n)$ with welfare OPT. As a warm up, observe that since the groups of buyers and sellers are chosen uniformly at random, the welfare maximizing allocation restricted to agents in B, assuming all items are available, has welfare of at least $E[OPT/2]$.

We now compute the optimal welfare with only agents in B assuming only items from $E_{\hat{S}}$ and that each buyer in B keeps his endowed items. Note that for every item j, $\Pr[j \in E_{\hat{S}}] = \frac{1}{4}$. To see this, consider agent i for which $j \in E_i$. Agent i is assigned to the sellers group S with exactly probability $\frac{1}{2}$ and agrees to sell with (independent) probability $\frac{1}{2}$ since he is offered his median price.

Fix some buyer $i \in B$: for each item $j \in O_i$, either $j \in E_i$ or $Pr[j \in E_{\hat{S}}] = \frac{1}{4}$. Thus, by Claim 3.2, $E[v_i(E_i + O_i \cap E_{\hat{S}})] \geq \frac{v_i(O_i)}{4}$ (we invoke the claim since E_i and the items in O_i are added to the bundle with probability of at least $\frac{1}{4}$). By linearity of expectation and recalling that an agent is in B with probability exactly $\frac{1}{2}$, the optimal welfare with only agents in B assuming only items from $E_{\hat{S}}$ and that each buyer in B keeps his endowed items is at least $E[OPT/8]$.

To compute the expected welfare of the allocation of the algorithm, we separately compute the expected welfare of the buyers and of the sellers. The welfare of the buyers is the optimal allocation among agents in B where each agent keeps his endowment and allocation A is penalized by c_A. Denote by Q the sum of the medians of players that agreed to sell, and observe that for every A, $c_A \leq H_t \cdot Q$. We will consider the allocation A' where each buyer $i \in B$ is assigned $E_i + O_i \cap E_{\hat{S}}$. This allocation is in the range of the VCG mechanism, thus the welfare of the allocation that VCG outputs is in expectation at least $\max(0, E[OPT/8] - E[H_t \cdot Q])$ (because the welfare of the allocation where no items is allocated is 0).

We now compute the expected welfare of the sellers. In fact, we take into account the welfare of sellers who declined to sell their items. Since by definition the value $v_i(E_i)$ of every seller i that declined to sell is above his median, we have that the expected welfare of sellers that declined to sell their items is at least $E[Q]$.

We have that the expected welfare of the allocation that the mechanism outputs is at least $\max(0, E[OPT/8] - E[H_t \cdot Q]) + E[Q]$. If $E[Q] \geq \frac{E[OPT]}{8H_t}$ then we are guaranteed an approximation ratio of at least $8H_t$. Else, we get that

$$
\begin{aligned}
&E[OPT/8] - E[H_t \cdot Q] + E[Q] &(1)\\
&= \frac{1}{8}E[OPT] - (H_t - 1)E[Q] &(2)\\
&> \frac{1}{8}E[OPT] - (H_t - 1)\frac{E[OPT]}{8H_t} = \frac{E[OPT]}{8H_t}, &(3)
\end{aligned}
$$

proving the desired approximation ratio. □

We note that although Theorem 3.1 is stated for the general model of combinatorial exchanges, the same mechanism can clearly be used for two-sided markets, i.e., where agents are divided in advance to roles of sellers and buyers. In this case, one should just skip the first stage that randomly assigns agents to roles, and thus improve the approximation ratio by a factor of 2. A recent paper by Colini-Baldeschi et al. [9] showed how to transform the Combinatorial Median Mechanism (for 2-sided markets) into a *strongly* budget-balanced mechanism by giving all the leftover money to a randomly chosen agent. This proved the existence of a strongly budget balanced, $4H_t$-approximation mechanism for 2-sided combinatorial auctions with sub-additive agents.

4 Arrow-Debreu Markets

In this section we give a constant approximation mechanism for a multi-parameter environment without any distributional assumptions. We consider the following n-agent exchange setting with a divisible good (the same model was studied in the classic work of Arrow and Debreu [2]): there is one divisible good and n players. Each player i has a valuation function $v_i : [0, 1] \to \mathbb{R}$, and for every x, y we define the marginal valuation $v_i(x|y) = v_i(x+y) - v_i(y)$. We assume that the valuation functions are normalized ($v_i(0) = 0$), non decreasing, and have decreasing marginal valuations (i.e., $v_i(\epsilon|x) \geq v_i(\epsilon|y)$ for every $\epsilon > 0, y > x$).[7] Denote the initial endowment of player i by r_i, $r_i \geq 0$, where $\Sigma_i r_i = 1$. Observe that given some price p, the supply that a seller i is willing to sell is an amount x_i that maximizes his payoff $p \cdot x_i - v_i(x_i | r_i - x_i)$.

[7] When $v_i(\cdot)$ is twice differentiable, we simply assume that $v_i^{''}(x) \leq 0$ for every x.

Our starting point is the combinatorial median mechanism for combinatorial exchanges. Similarly, players are divided into "sellers" and "buyers" (but in a subtler way). The constant approximation ratio is achieved by replacing the revenue extraction procedure of the combinatorial-exchange mechanism with a method that allows the separate sale of items. This will be the key to obtaining a constant approximation in the worst case. We assume that no single player initially holds a huge chunk of the good; Notice that if, say, one player is initially endowed with all the good then no dominant-strategy prior-free mechanism can achieve a bounded approximation (such mechanisms essentially post a price to the agents, and this price may not clear the market, see [21]).

Theorem 4.1. *Suppose that for all i, $r_i \leq \frac{1}{8}$. Then, there exists a truthful, weakly budget balanced mechanism that provides an expected approximation ratio of 48 in every instance. The theorem holds for any number of agents and for every profile of preferences satisfying decreasing marginal valuations.*

We first provide a mechanism assuming that the bidders can be divided to 3 groups N_1, N_2, N_3 where each set N_k is *substantial*: $\Sigma_{i \in N_k} r_i \geq \frac{1}{4}$. Later we relax the requirement; we will only assume that for every i, $r_i \leq \frac{1}{8}$. We need the following definition:

Definition 4.2. *Let N_k be a substantial set of bidders. The* mid-supply price *of N_k is the minimal price p such that the total amount that bidders in N_k are willing to sell at price p is at least $\frac{1}{8}$.*

The mechanism itself is a bit heavy on details, although the basic idea is quite simple. We therefore start with an informal description, and then move on to a formal one. Initially, we have an arbitrary division of the bidders into three substantial groups. Select their roles at random so N_1 is the group of buyers, N_2 is a set of players that will provide us with statistics, and let N_3 be the set of sellers.

We use the statistics group to compute a mid-supply price p: that is, the price for which agents in the statistics group are willing to sell half of their total supply (which is at least $\frac{1}{8}$ of the good). Each seller in N_3 is asked to report the maximum amount of his endowment that he is willing to sell at price p. Let t be the total amount that the sellers are willing to sell at price p (in fact, we have to make sure that $t \leq \frac{1}{8}$ – see the formal description for exact implementation details).

Now, we run the VCG mechanism with the participation of the buyers and an extra buyer with valuation $v_d(s) = \min\{t, s\} \cdot p$. This bidder is added to ensure compliance with the budget balance requirement. The set of possible allocations in this VCG mechanism is all allocations of amount t of the good among the buyers N_1 and the extra buyer. Effectively, we show that this amounts to finding a welfare-maximizing allocation of t fraction of the good among the buyers so that each buyer that received an amount of x pays at least $x \cdot p$. We use this money to pay each seller $i \in N_3$ a total sum of $x_i \cdot p$, where x_i is the part of the endowment that was taken from bidder i. We now provide a formal description of the mechanism, followed by its analysis.

The Formal Mechanism:

1. Divide the agents into three substantial groups denoted by N_1, N_2 and N_3 based on their endowments r_i.[8]

[8] To accomplish that, greedily add agents to N_1 while the sum of the r_i's of agents in N_1 is at most $\frac{1}{8}$. Stop adding agents to N_1 when we add an agent that makes an "overflow": $\Sigma_{i \in N_1} r_i > \frac{1}{8}$. Since each $r_i \leq \frac{1}{8}$, we also have that $\Sigma_{i \in N_1} r_i \leq \frac{1}{4}$. Continue similarly, only with agents that were not added to N_1, to construct N_2 and N_3.

2. Select uniformly at random "roles" for the groups N_1, N_2, N_3: players in one group will be the *buyers* (without loss of generality, N_1), another group will be the *statistics* group (without loss of generality, N_2), and players in the additional group are the *sellers* (N_3).

3. Agents in N_2 report their preferences to the mechanism, and let p be the mid-supply price of the statistics group N_2. Each seller $i \in N_3$ reports the amount of good x'_i he is willing to sell at price p. Let $t = \min\{\frac{1}{8}, \Sigma_{i\in N_3} x'_i\}$.

   If $\Sigma_{i\in N_3} x'_i \leq \frac{1}{8}$, let $x_i = x'_i$. If $\Sigma_{i\in N_3} x'_i > \frac{1}{8}$, choose a value x_i for each i such that $x_i \leq x'_i$ and $\Sigma_{i\in N_3} x_i = \frac{1}{8}$.[9]

4. Run the VCG mechanism to sell t-fraction of the good to the buyers in N_1 and to a dummy buyer, where the valuation of each buyer $i \in N_1$ is $v'_i = v_i(s|r_i)$ and the valuation of the dummy buyer is $v_d(s) = \min\{t, s\} \cdot p$.

5. The output of the mechanism is as follows: each buyer $i \in N_1$ pays to the mechanism the VCG payment of v'_i and receives the same amount of good that v'_i received (in addition to his endowment r_i). Agents in N_2 keep their initial endowment and do not pay anything.

   Let $t' \leq t$ be the amount of good that the dummy player ended up with in the VCG mechanism. Choose x''_i's such that for each $i \in N_3$, $x''_i \leq x_i$ and $\Sigma_{i\in N_3} x''_i = t - t'$, taking the good first from sellers with lower indices. Each seller $i \in N_3$ keeps $r_i - x''_i$ of the good and receives a payment of $x''_i \cdot p$.

Claim 4.3. *The above mechanism is truthful.*

Proof. The mechanism is clearly truthful for the statistics group since they never sell nor receive any amount of the good. The mechanism is also truthful for the buyers since they are just participating in a VCG mechanism. To show that the mechanism is truthful for the sellers we have to use the fact that the valuations exhibit decreasing marginals. First, observe that the price p depends only on the valuations of the statistics group N_2. Now, consider some seller $i \in N_3$. Seller i reports at stage 3 a quantity x'_i that maximizes his profit $x \cdot p + v_i(r_i - x)$. Therefore, if she eventually sells a fraction x'_i she has no reason to report a different value. However, at stages 3 and 5 of the mechanism the quantity that she sells is reduced to x_i or x''_i that may gain her a lower profit. By the way that the mechanism reduces the quantities, reporting any value above x''_i will not affect the quantity that seller i sells. If seller i reports a value smaller than x''_i, he will sell a quantity smaller than x''_i; This smaller quantity cannot gain him a greater profit since the profit is non-decreasing in x in the range below x'_i due to decreasing marginals.[10] Finally, we assign agents to the groups N_1, N_2 and N_3 based only on their shares r_i which is public knowledge. □

Claim 4.4. *The above mechanism is weakly budget balanced.*

Proof. Consider a buyer i that received an amount of t_i (not including his endowment r_i). His VCG price is at least $t_i \cdot p$, since otherwise we could have considered the same allocation except that an additional amount of t_i of the good is allocated to the dummy bidder. We have that the total payment is at least $(t - t') \cdot p$, which is exactly the amount we have to pay to the sellers. □

The next lemma analyses the approximation ratio of the mechanism:

[9] Formally, order the buyers arbitrarily, and let $x_i = \max\{0, \frac{1}{8} - \Sigma_{i'>i} x'_{i'}\}$ if $x'_i + \Sigma_{i'>i} x'_{i'} \geq \frac{1}{8}$.

[10] To see this, note that the derivative of $x \cdot p$ is p for every x. Since x'_i maximizes profit, and due to the convexity of v_i, for every value $x < x'_i$ the derivative of $v_i(r_i - x)$ is negative with absolute value of at most p. Therefore, the marginal profit is non-negative for $x < x'_i$. A similar argument holds also when v_i is not differentiable.

Lemma 4.5. *The mechanism provides an approximation ratio of* 48 *for any profile of agents with decreasing marginal valuations.*

Proof. Fix some optimal solution to the original exchange problem $(o_1, \ldots, o_n)$. For $k \in \{1, 2, 3\}$, let $O_k = \Sigma_{i \in N_k} v_i(o_i)$. Observe that since each group of bidders plays the role of the buyers with probability exactly $\frac{1}{3}$, we have that $E[O_1] = OPT/3$. Let p' be the mid-supply price of N_3 and recall that p is the mid-supply price of N_2. Since the statistics group and the sellers group are chosen at random, with probability at least $1/2$ we have that $p \geq p'$. We will condition our analysis on that event and conservatively assume that if $p < p'$ then the welfare of the allocation that the mechanism outputs is 0.

Now, if $p \geq p'$, the total amount of the good that bidders in N_3 are willing to sell is at least $\frac{1}{8}$: they hold at least $\frac{1}{4}$ of the good since N_3 is "substantial", and at price p' they are willing to sell half of it, so surely they will agree to sell that amount at price $p \geq p'$. In particular, we have that $\Sigma_i x_i = \frac{1}{8}$ (and thus $t = 1/8$) with probability at least $\frac{1}{2}$.

Claim 4.6. *For every $i \in N_1$, let s_i denote the amount of good bidder i receives in the final allocation of our mechanism. If $t = \frac{1}{8}$ then $\Sigma_{i \in N_1} v_i(s_i) \geq \frac{O_1}{8} - \frac{p}{8}$.*

Proof. Consider the allocation that gives each bidder $i \in N_1$ an amount of $s'_i = \frac{o_i}{8}$. Since the valuations have decreasing marginals, the welfare of this allocation is at least $\Sigma_i v_{i \in N_1}(s'_i) \geq \frac{O_1}{8}$ and no more than $\frac{1}{8}$ of the good was allocated to players in N_1.

Thus, the welfare of VCG is at least the welfare of $(s'_1, \ldots, s'_n)$, but we have to subtract the contribution of the dummy buyer. His contribution to the welfare is at most his maximum value (obtained when he is allocated all the $t = 1/8$): $\frac{p}{8}$. Thus, we have that $\Sigma_{i \in N_1} v_i(s'_i) \geq \frac{O_1}{8} - \frac{p}{8}$. □

Now notice that the value of bidders in the statistics group N_2 is at least $p \cdot t = \frac{p}{8}$: they are not willing to sell $\frac{1}{8} = t$ of the good at price p, so their total value for their initial endowment (that they keep) is at least $p \cdot t$.

Hence, we have that with probability at least $\frac{1}{2}$ (if $t = \frac{1}{8}$), it holds that $\Sigma_{i \in N_1} v_i(s_i) + \Sigma_{i \in N_2} v_i(s_i) = \frac{O_1}{8} - \frac{p}{8} + \frac{p}{8} = \frac{O_1}{8}$ (where $s_1, .., s_n$ is the allocation by the mechanism). Recall that $E[O_1] = \frac{OPT}{3}$, and we get that the expected welfare is at least $\frac{1}{2} \cdot \frac{E[O_1]}{8} \geq \frac{OPT}{48}$, as needed. □

We note that although we discuss a model with a divisible good, our result also applies for reallocating units of a discrete homogenous good among agents. As long as each agent initially holds at most 1/8 of the units, and that all preferences satisfy decreasing marginal utilities, Theorem 4.1 should hold (up to minor rounding issues).

References

[1] George A Akerlof. The market for 'lemons': Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.

[2] Kenneth J. Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):pp. 265–290, 1954.

[3] Lawrence Ausubel, Jon Levin, Paul Milgrom, and Ilya Segal. Incentive auction rules option and discussion, 2013. Appendix to the FCCs 28-Sep-2012 NPRM on Incentive Auctions.

[4] Moshe Babaioff, Patrick Briest, and Piotr Krysta. On the approximability of combinatorial exchange problems. In *SAGT*, pages 83–94, 2008.

[5] Liad Blumrosen and Shahar Dobzinski. (almost) efficient mechanisms for bilateral trading. *CoRR, Working paper.*, abs/1604.04876, 2016.

[6] Liad Blumrosen and Yehonatan Mizrahi. Approximating gains-from-trade in bilateral trading. In *Web and Internet Economics WINE*, pages 400–413, 2016.

[7] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 589–590, 2017.

[8] Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *SODA 2016*, 2016.

[9] Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. Approximately efficient two-sided combinatorial auctions. In *The ACM Conference on Economics and Computation*, pages 591–608, 2017.

[10] Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Fixed price approximability of the optimal gain from trade. In *Web and Internet Economics - WINE*, pages 146–160, 2017.

[11] Peter Cramton, Robert Gibbons, and Paul Klemperer. Dissolving a partnership efficiently. *Econometrica*, 55(3):615–32, 1987.

[12] Martin W. Cripps and Jeroen M. Swinkels. Efficiency of large double auctions. *Econometrica*, 74(1):47–92, 2006.

[13] Shahar Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, STOC '11, pages 139–148, 2011.

[14] Shahar Dobzinski. Computational efficiency requires simple taxation. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 209–218, 2016.

[15] Shahar Dobzinski, Aranyak Mehta, Tim Roughgarden, and Mukund Sundararajan. Is shapley cost sharing optimal? In *International Symposium on Algorithmic Game Theory*, pages 327–336. Springer, 2008.

[16] Shahar Dobzinski and Shahar Ovadia. Combinatorial cost sharing. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 387–404. ACM, 2017.

[17] Shahar Dobzinski and Jan Vondrák. The computational complexity of truthfulness in combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 405–422, 2012.

[18] Shaddin Dughmi and Jan Vondrák. Limitations of randomized mechanisms for combinatorial auctions. In *FOCS*, pages 502–511, 2011.

[19] Paul Dutting, Tim Roughgarden, and Inbal Talgam-Cohen. Modularity and greed in double auctions. *EC'14*.

[20] Drew Fudenberg, Markus Mobius, and Adam Szeidl. Existence of equilibrium in large double auctions. *Journal of Economic Theory*, 133(1):550–567, 2007.

[21] Kathleen M Hagerty and William P Rogerson. Robust trading mechanisms. *Journal of Economic Theory*, 42(1):94 – 107, 1987.

[22] A. Kothari, T. Sandholm, and S. Suri. Solving combinatorial exchanges: optimality via a few partial bids. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004.*, pages 1418–1419, 2004.

[23] Paul Milgrom and Ilya Segal. Deferred-acceptance heuristic auctions, 2013. Working paper, Stanford University.

[24] Hervé Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.

[25] R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265 – 281, 1983.

[26] Noam Nisan. *In Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay Vazirani (Editors), Algorithmic Game Theory. Chapter 9. Introduction to Mechanism Design (for Computer Scientists).* Cambridge University Press., 2007.

[27] Tim Roughgarden and Mukund Sundararajan. New trade-offs in cost-sharing mechanisms. In *STOC*, pages 79–88, 2006.

[28] Aldo Rustichini, Mark A Satterthwaite, and Steven R Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5):1041–63, 1994.

[29] Mark A. Satterthwaite and Steven R Williams. The optimality of a simple market mechanism. *Econometrica*, 70(5):1841–1863, 2002.

[30] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. Sbba: A strongly-budget-balanced double-auction mechanism. In *the 9th International Symposium on Algorithmic Game Theory*, pages 260–272, 2016.